arXiv:1612.01829v1 [cs.DS] 6 Dec 2016

# Symmetry exploitation for Online Machine Covering with Bounded Migration*

Waldo Gálvez† José A. Soto‡ José Verschae§

December 7, 2016

**Abstract**

Online models that allow recourse are highly effective in situations where classical models are too pessimistic. One such problem is the online machine covering problem on identical machines. In this setting jobs arrive one by one and must be assigned to machines with the objective of maximizing the minimum machine load. When a job arrives, we are allowed to reassign some jobs as long as their total size is (at most) proportional to the processing time of the arriving job. The proportionality constant is called the *migration factor* of the algorithm.

Using a new rounding procedure specially tailored for online problems, we design a $(4/3+\varepsilon)$-competitive algorithm using migration factor $\tilde{O}(1/\varepsilon^3)$. At every arrival we run an adaptation of the *Largest Processing Time first* (LPT) algorithm. Since the new job can cause a complete change of the assignment of smaller jobs, a low migration factor is achieved by carefully exploiting the highly symmetric structure obtained by our rounding.

We also study local search algorithms for the machine covering problem, and show that *jump* and *swap* optimality have an approximation ratio which lies in the interval $[1.691, 1.75]$ and can be adapted to the online context with a small constant as migration factor. Our lower bound is obtained by a nice construction based on Sylvester's sequence.

## 1 Introduction

We consider a fundamental load balancing problem where $n$ jobs need to be assigned to $m$ identical parallel machines. Each job $j$ is fully characterized by its non-negative processing time $p_j$. Given an assignment of jobs, the load of a machine is the sum of the processing times of jobs assigned to it. The *machine covering problem* asks for an assignment of jobs to machines maximizing the load of the least loaded machine.

This problem is well known to be strongly NP-hard and allows for a polynomial-time approximation scheme (PTAS) [21]. A well studied algorithm for this problem is the *Largest Processing Time First* rule (LPT), that sorts the jobs non-increasingly and assigns them iteratively to the least

*This work was partially supported by FONDECYT project 11140579, FONDECYT project 11130266, and by Nucleo Milenio Información y Coordinación en Redes ICM/FIC RC130003.

†IDSIA, USI-SUPSI, Lugano, Switzerland. E-mail: waldo@idsia.ch

‡Departamento de Ingeniería Matemática & CMM, Universidad de Chile, Santiago, Chile. E-mail: jsoto@dim.uchile.cl

§Facultad de Matemáticas & Escuela de Ingeniería, Pontificia Universidad Católica de Chile, Santiago, Chile. E-mail: jverschae@uc.cl

loaded machine. Deuermeyer et al. [6] show that LPT is a $\frac{4}{3}$-approximation and that this factor is asymptotically tight; later Csirik et al. [5] refine the analysis giving a tight bound for each $m$.

In the online setting jobs arrive one after another, and at the moment of an arrival we must decide on a machine to assign the arriving job. This natural problem does not admit a constant competitive ratio. Deterministically, the best possible competitive ratio is $m$ [21], while randomization allows for a $\tilde{O}(\sqrt{m})$-competitive algorithm, which is the best possible up to logarithmic factors [1].

**Dynamic model.** The previous negative facts motivate the study of a relaxed online scenario with *bounded migration*. Unlike the classic online model, when a new job $j$ arrives we are allowed to reassign other jobs. More precisely, given a constant $\beta > 0$, we can migrate jobs whose total size is upper bounded by $\beta p_j$. The value $\beta$ is called the *migration factor* and it accounts for the robustness of the algorithm. In one extreme, we can model the usual online framework by setting $\beta = 0$. In the other extreme, setting $\beta = \infty$ allows to compute the optimal offline solution in each iteration. Our main interest is to understand the exact trade-off between the migration factor $\beta$ and the competitiveness of our algorithms. Besides being a natural problem with an interesting theoretical motivation, its original purpose was to find good algorithms for a problem in the context of Storage Area Networks (SAN) [18].

**Local search and migration.** The local search method has been extensively used to tackle different hard combinatorial problems. Local search algorithms and online algorithms where recourse is allowed are closely related. This comes from the fact that simple local search neighborhoods allow to get considerably improved solutions while having accurate control over the recourse actions needed, and in some cases even a bounded number of local moves leads to substantially improved solutions (see [10, 13, 15] for examples in network design problems).

**Related Work.** Sanders et al. [18] develop online algorithms for load balancing problems in the migration framework. For the makespan minimization objective, where the aim is to minimize the maximum load, they give a $(1+\varepsilon)$-competitive algorithm with migration factor $2^{\tilde{O}(1/\varepsilon)}$. A mayor open problem in this area is to determine whether a migration factor of $\text{poly}(1/\varepsilon)$ is achievable.

The landscape for the machine covering problem is somewhat different. Sanders et al. [18] give a 2-competitive algorithm with migration factor 1, this is until now the best competitive ratio known for any algorithm with constant migration factor. On the negative side, Skutella and Verschae [20] show that it is not possible to maintain arbitrarily near optimal solutions using constant migration factor, giving a lower bound of 20/19 for the best competitive ratio achievable in that case. The lower bound is based on an instance where arriving jobs are very small, which do not allow to migrate any other job. This motivated the study of an amortized version, called *reassignment cost model*, and develop a $(1+\varepsilon)$-competitive algorithm using constant reassignment factor. They also show that if all arriving jobs are larger that $\varepsilon \cdot \text{OPT}$, then there is a $(1+\varepsilon)$-competitive algorithm with constant migration factor.

Similar migration models have been studied for other packing and covering problems. For example Epstein & Levin [7] design a $(1+\varepsilon)$-competitive algorithm for online bin packing using $2^{\tilde{O}(1/\varepsilon^2)}$ migration factor , which was improved later by Jansen & Klein [11] to $\text{poly}(1/\varepsilon)$ migration factor, and then further refined by Berndt et al. [3]. Also, for makespan minimization with

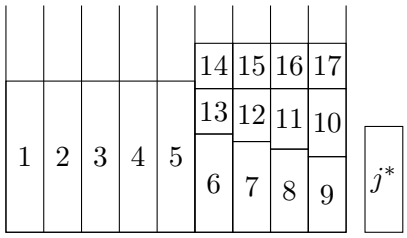


(a) LPT for the original instance and arriving job $j^*$.

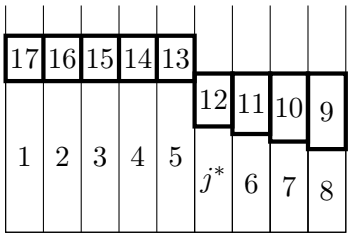


(b) LPT for the new instance. Thick items correspond to migrated jobs.

Figure 1: $\Omega(m)$ migration factor needed to maintain LPT at the arrival of $j^*$.

preemption and other objectives, Epstein & Levin [8] design a best-possible online algorithm using migration factor $\left(1 - \frac{1}{m}\right)$.

Regarding local search applied to load balancing problems, many neighborhoods have been studied such as *jump*, *swap*, *push* and *lexicographical jump* in the context of makespan minimization on related machines [19], makespan minimization on restricted related machines [16], and also multi-exchange neighborhoods for makespan minimization on identical parallel machines [9]. To the best of our knowledge, this is the first time that local search algorithms are analyzed for the machine covering problem.

**Our Contribution.** Our main result is a $(4/3+\varepsilon)$-competitive algorithm using poly$(1/\varepsilon)$ migration factor. This is achieved by running a carefully crafted version of LPT at the arrival of each new job. We would like to stress that, even though LPT is a simple and very well studied algorithm in the offline context, solutions can change dramatically in the online context; see Figure 1 and Lemma 28 in Appendix A.

To overcome this barrier, we first adapt a less standard rounding procedure to the online framework. The rounding is such that jobs larger than $\Omega(\varepsilon\text{OPT})$ (where OPT is the offline optimum value) take only $\tilde{O}(1/\varepsilon)$ many job sizes, and all values are multiples of a common number $g \in \Theta(\varepsilon^2\text{OPT})$. This implies that the number of possible loads for machines having only big jobs is constant. Unlike known techniques that yield similar results (see e.g. [12]), our rounding is well suited for online algorithms and helps simplifying the analysis as it does not depend on OPT (which varies through iterations).

The properties obtained by rounding are crucial to create symmetries: While assigning a job to a current least loaded machine, since there is a constant number of possible machine loads, there will usually be multiple least loaded machines to assign the job. All options lead to different solutions in terms of job assignments, but they all achieve the same objective value. Broadly speaking, the algorithm will construct one of these equivalent schedules, trying to maintain as many machines with assignments equivalent to the original solution as possible. The analysis of the algorithm will rely on monotonicity properties implied by LPT which, coupled with rounding, implies that for every job size the increase in the number of machines with different assignments (w.r.t the solution of the previous iteration) is constant, finally yielding a migration factor that only grows polynomially on $1/\varepsilon$. In addition, we give a slightly improved lower bound of 17/16 for the best competitive ratio achievable using constant migration.

In the second part of the paper we analyze the local search neighborhoods jump and swap. We show that for the machine covering problem they give an approximation guarantee $r$ lying in

the interval $[1.691, 1.75]$. The family of instances leading to the lower bound is constructed using Sylvester's sequence and exploiting its properties. Moreover, using the rounding procedure for the online setting, the local search procedures can be adapted to the online setting. Compared to our online version of LPT, this yields a slightly worse competitive factor of $(r+\varepsilon)$ but gets a improved migration factor of $O(1/\varepsilon)$.

## 2 Robust LPT with migration factor poly$(1/\varepsilon)$.

Consider a set of $n$ jobs $\mathcal{J}$ and a set of $m$ machines $\mathcal{M}$. In our problem, a solution or schedule $\mathcal{S} : \mathcal{J} \to \mathcal{M}$ corresponds to an assignment of jobs to machines. The set of jobs assigned to a machine $i$ is then $\mathcal{S}^{-1}(i) \subseteq J$. The load of machine $i$ in $\mathcal{S}$ corresponds to $\ell_i(\mathcal{S}) = \sum_{j \in \mathcal{S}^{-1}(i)} p_j$. The minimum load is denoted by $\ell_{\min}(\mathcal{S}) = \min_{i \in \mathcal{M}} \ell_i(\mathcal{S})$, and a machine $i$ is said to be *least loaded* in $\mathcal{S}$ if $\ell_i(\mathcal{S}) = \ell_{\min}(\mathcal{S})$.

For an algorithm $\mathcal{A}$ and set of jobs $\mathcal{J}$ we denote by $\mathcal{S}_{\mathcal{A}}(\mathcal{J}, \mathcal{M})$ the schedule returned by $\mathcal{A}$ on instance $(\mathcal{J}, \mathcal{M})$. Similarly, $\mathcal{S}_{\mathrm{OPT}}(\mathcal{J}, \mathcal{M})$ denotes the optimal schedule, being $\mathrm{OPT}(\mathcal{J}, \mathcal{M})$ its minimum load. When it is clear from the context we will drop the dependency on $\mathcal{J}$ or $\mathcal{M}$.

### 2.1 Algorithms with robust structure.

An important fact used in the design of the robust PTAS for makespan minimization from Sanders et al. [18] is that small jobs can be assigned greedily almost without affecting the approximation guarantee. This is however not the case for machine covering; see, e.g. [20] or Lemma 19. One way to avoid this inconvenience is to develop algorithms that are oblivious to the arrival of small jobs, that is, algorithms where the assignment of big jobs is not affected when a new small job arrives.

**Definition 1.** *Let $h \in \mathbb{R}_+$. An algorithm $\mathcal{A}$ has **robust structure at level** $h$ if, for any instance $(\mathcal{J}, \mathcal{M})$ and $j^* \notin \mathcal{J}$ such that $p_{j^*} \leq h$, $\mathcal{S}_{\mathcal{A}}(\mathcal{J}, \mathcal{M})$ and $\mathcal{S}_{\mathcal{A}}(\mathcal{J} \cup \{j^*\}, \mathcal{M})$ assign to the same machines all the jobs in $\mathcal{J}$ with processing time at least $h$.*

This definition highlights also the usefulness of working with the LPT rule, since the addition of a new small job to the instance does not affect the assignment of larger jobs. Indeed, it is easy to see the following.

**Remark.** *For any $h \in \mathbb{R}_+$,* LPT *has robust structure at level $h$.*

The next definition will help show that algorithms with robust structure at level $h$, for small enough $h$, allow us to focus mainly on maintaining the structure of the big jobs.

**Definition 2.** *Let $\mathcal{A}$ be an $\alpha$-approximation algorithm for the machine covering problem, with $\alpha$ constant, $k \in \mathbb{N}$ constant, and $\varepsilon > 0$. Given a machine covering instance $(\mathcal{J}, \mathcal{M})$, a schedule $\mathcal{S}$ is a **$k$-relaxed version of** $\mathcal{S}_{\mathcal{A}}$ if:*

1. *jobs with processing time at least $k\varepsilon\mathrm{OPT}$ are assigned exactly as in $\mathcal{S}_{\mathcal{A}}$, and*
2. *for every machine $i \in \mathcal{M}$ such that $\mathcal{S}$ assigns at least one job of size less than $k\varepsilon\mathrm{OPT}$, $\ell_i(\mathcal{S}) \leq \ell_{\min}(\mathcal{S}) + k\varepsilon\mathrm{OPT}$.*

**Lemma 3.** *Let $\mathcal{A}$ be an $\alpha$-approximation, $\alpha \geq 1$ constant, $k \in \mathbb{N}$ constant, $0 < \varepsilon < \frac{1}{2k\alpha}$ and $(\mathcal{J}, \mathcal{M})$ a machine covering instance. If $\mathcal{A}$ has robust structure at level $k\varepsilon\mathrm{OPT}$, then every $k$-relaxed version of $\mathcal{S}_{\mathcal{A}}$ is an $(\alpha + O(\varepsilon))$-approximate solution.*

*Proof.* Suppose by contradiction that there exists a $k$-relaxed version of $\mathcal{S}_{\mathcal{A}}$, say $\mathcal{S}$, which is not $(\alpha + 2k\alpha^2\varepsilon)$-approximate. This implies that $\ell_{\min}(\mathcal{S}) < \frac{1}{\alpha+2k\alpha^2\varepsilon}\text{OPT} \leq \left(\frac{1}{\alpha} - k\varepsilon\right)\text{OPT}$.

Let $\mathcal{M}_s$ the set of machines where $\mathcal{S}$ assigns at least one job of size less than $k\varepsilon\text{OPT}$. Notice that $\mathcal{M}_s \neq \emptyset$, because otherwise $\ell_{\min}(\mathcal{S}_{\mathcal{A}}) = \ell_{\min}(\mathcal{S}) < \left(\frac{1}{\alpha} - k\varepsilon\right)\text{OPT}$, which contradicts that $\mathcal{S}_{\mathcal{A}}$ is $\alpha$-approximate. Since $\mathcal{A}$ has robust structure at level $k\varepsilon\text{OPT}$, $\mathcal{S}$ and $\mathcal{S}_{\mathcal{A}}$ differ only on the machines from $\mathcal{M}_s$, and the total processing time of jobs assigned by $\mathcal{S}$ to $\mathcal{M}_s$ is at most $|\mathcal{M}_s|(\ell_{\min}(\mathcal{S}) + k\varepsilon\text{OPT})$. Thus,

$$\ell_{\min}(\mathcal{S}_{\mathcal{A}}) \leq \min_{i\in\mathcal{M}_s} \ell_i(\mathcal{S}_{\mathcal{A}}) \leq \ell_{\min}(\mathcal{S}) + k\varepsilon\text{OPT} < \frac{1}{\alpha}\text{OPT},$$

which contradicts that $\mathcal{S}_{\mathcal{A}}$ is $\alpha$-approximate. □

Assume that we have developed an algorithm $\mathcal{A}$ with robust structure to some convenient level $h$, such that the migration of jobs larger than $h$ is upper bounded properly. Dealing with jobs smaller than $h$ becomes simpler since, as we will see, a simple greedy algorithm helps to assign small jobs, creating a $k$-relaxed version of $\mathcal{S}_{\mathcal{A}}$.

## 2.2 Rounding, LPT and symmetry exploitation

In this section we show first how to create symmetries via rounding the job sizes and how to exploit this to obtain bounded migration factors.

**Rounding.** Let us consider $0 < \varepsilon < 1$ such that $1/\varepsilon \in \mathbb{Z}$. We introduce a new rounding technique that is specially designed for our online algorithm. For any job $j$, let $e_j \in \mathbb{Z}$ be such that $2^{e_j} \leq p_j < 2^{e_j+1}$. We then round down $p_j$ to the previous number of the form $2^{e_j} + k\varepsilon 2^{e_j}$ for $k \in \mathbb{N}$, that is, we define $\tilde{p}_j := 2^{e_j} + \left\lfloor \frac{p_j - 2^{e_j}}{\varepsilon 2^{e_j}} \right\rfloor \varepsilon 2^{e_j}$.

Observe that $p_j \geq \tilde{p}_j \geq p_j - \varepsilon 2^{e_j} \geq (1-\varepsilon)p_j$. Hence, an $\alpha$-approximation algorithm for a rounded instance is immediately an $(\alpha/(1-\varepsilon))$-approximation for the original instance. From now on we work exclusively with the rounded processing times.

Consider an upper bound UB on OPT such that $\text{OPT} \leq \text{UB} \leq 2\text{OPT}$. This can be computed with any 2-approximation to the problem, in particular from LPT. Consider the index set

$$\tilde{I}(\text{UB}) := \left\{i \in \mathbb{Z} : \varepsilon\text{UB} \leq 2^i < \text{UB}\right\} = \{\ell, \ldots, u\}. \tag{1}$$

We classify jobs as *small* if $\tilde{p}_j < 2^\ell$, *big* if $\tilde{p}_j \in [2^\ell, 2^{u+1})$, and *huge* otherwise. Notice that small jobs have size at most $2\varepsilon\text{UB}$ and huge jobs have size at least UB. As we will see, our main difficulty will be given by big jobs; small and huge jobs are easy to handle. Notice that in every solution $\mathcal{S}$ constructed using LPT, without considering small jobs, huge jobs are assigned to a machine on their own and every machine without huge jobs has load at most $2\text{UB}$ (either it has a big job alone, of size at most $2\text{UB}$, or it has load at most $\ell_{\min}(\mathcal{S}) + \tilde{p}_j \leq 2\ell_{\min}(\mathcal{S}) \leq 2\text{UB}$, where $j$ is the smallest job in the machine). Let

$$\tilde{P} = \left\{2^i + k\varepsilon 2^i : i \in \{\ell, \ldots, u\}, k \in \{0, 1, \ldots, (1/\varepsilon) - 1\}\right\}, \tag{2}$$

the set of all (rounded) processing times that a big job may take. The next lemma highlights the main properties of our rounding procedure.

**Lemma 4.** *Consider the rounded job sizes $\tilde{p}_j$ for all $j$. Then it holds that,*

1. $|\tilde{P}| \in O((1/\varepsilon)\log(1/\varepsilon))$, *and*
2. *for each big and huge job $j$ it holds that $\tilde{p}_j = h \cdot \varepsilon 2^\ell$ for some $h \in \mathbb{N}_0$.*

*Proof.* From the definition of $\tilde{P}$, we have that $|\tilde{P}| = \frac{1}{\varepsilon}(u - \ell - 1) - 1$. Since $\ell \geq \log(\varepsilon \text{UB})$ and $u \leq \log(\text{UB})$, then $(u - \ell - 1) \leq \log\left(\frac{1}{\varepsilon}\right)$. Altogether, $|\tilde{P}| \in O((1/\varepsilon)\log(1/\varepsilon))$. Also, if $j$ is a big or huge job, then $\tilde{p}_j = 2^i + k\varepsilon 2^i$ for some $i \geq \ell$ and $k \in \{0, \ldots, \frac{1}{\varepsilon} - 1\}$. We conclude by noticing that $\tilde{p}_j = \left(\frac{1}{\varepsilon} + k\right) 2^{i-\ell} \cdot \varepsilon 2^\ell = h \cdot \varepsilon 2^\ell$ for $h = \left(\frac{1}{\varepsilon} + k\right) 2^{i-\ell} \in \mathbb{N}_0$. □

Unlike other standard techniques [12], the rounded sizes do not depend on OPT (or UB). This avoids possible migrations provoked by new rounded values, greatly simplifying our techniques.

**Load Monotonicity.** The algorithm described in the present section is an approximate online adaptation of LPT using poly$(1/\varepsilon)$ migration factor. In order to analyze it we will first show some structural properties of the solutions constructed by LPT and how it behaves when the instance is perturbed by a new job. Unless stated otherwise, for the rest of this section machine loads are considered with respect to the rounded processing times $\tilde{p}_j$.

**Definition 5.** *Given a schedule $\mathcal{S}$, its **load profile**, denoted by* $\text{load}(\mathcal{S})$*, is an $\mathbb{R}^m_{\geq 0}$-vector $(t_1, \ldots, t_m)$ containing the load of each machine sorted so that $t_1 \leq t_2 \leq \ldots \leq t_m$. .*

The following lemma shows that after the arrival of a job, the load profile of solutions constructed using LPT can only increase. We remark that this property only holds if the vector of loads is sorted, as it can be observed in Figure 1. This monotonicity will be crucial later for our analysis. To show the mentioned property, the following rather technical lemma will help.

**Lemma 6.** *Let $x, y \in \mathbb{R}^n_+$, $x = (x_1, x_2, \ldots, x_n)$, $y = (y_1, y_2, \ldots, y_n)$ such that $x_1 \leq x_2 \leq \cdots \leq x_n$, $y_1 \leq y_2 \leq \cdots \leq y_n$ and $x \leq y$ coordinate-wise, and $\alpha, \beta \in \mathbb{R}$ such that $\alpha \leq \beta$. If we consider the new vectors defined by replacing $x_i$ by $x_i + \alpha$ in $x$ and $y_i$ by $y_i + \beta$ in $y$ for some $i \in \{1, 2, \ldots, n\}$, and then we sort the coordinates non-decreasingly of the new vectors, obtaining $x'$ and $y'$, then $x' \leq y'$ coordinate-wise.*

*Proof.* Let $\bar{i}$ be the coordinate such that $x'_{\bar{i}} = x_i + \alpha$ and $\bar{j}$ such that $y'_{\bar{j}} = y_i + \beta$. For each coordinate $k < \min\{\bar{i}, \bar{j}\}$ or $k > \max\{\bar{i}, \bar{j}\}$ we have that $x'_k = x_k$ and $y'_k = y_k$, thus satisfying the desired inequality by hypothesis. In the remaining we have two cases:

- Assume that $\bar{i} < \bar{j}$, and let $k \in \{\bar{i}, \bar{i} + 1, \ldots, \bar{j} - 1\}$. Then we have that

  $$x'_k \leq x'_{k+1} = x_{k+1} \leq y_{k+1} = y'_k$$

  where the first inequality holds due to the monotonicity of the vectors and the second one because of the hypothesis. Similarly, for $k = \bar{j}$ we have that

  $$x'_k = x_k \leq y_k = y'_{k-1} \leq y'_k,$$

  where the first inequality follows from the hypothesis.

- Assume that $\bar{i} \geq \bar{j}$ and let $k \in \{\bar{j}, \bar{j}+1, \ldots, \bar{i}\}$. Then,

$$x'_k \leq x'_{\bar{i}} \leq y'_{\bar{j}} \leq y'_k,$$

where, the first and third inequalities follows from the monotonicity of the vectors, and the second one from the fact that $x_i + \alpha \leq y_i + \beta$. □

**Lemma 7.** *Let $(\mathcal{J}, \mathcal{M})$ be a machine covering instance and $j^* \notin \mathcal{J}$ a job. Then, it holds that* $\text{load}(\mathcal{S}_{\text{LPT}}(\mathcal{J}, \mathcal{M})) \leq \text{load}(\mathcal{S}_{\text{LPT}}(\mathcal{J}', \mathcal{M}))$, *where the inequality is considered coordinate-wise and $\mathcal{J}' = \mathcal{J} \cup \{j^*\}$.*

*Proof.* Let us first relabel the jobs in $\mathcal{J}$ so that $\tilde{p}_1 \geq \tilde{p}_2 \geq \ldots \geq \tilde{p}_n$. To simplify the argument we assume that both runs of LPT assign jobs in the order given by the labeling above $1, 2, \ldots, n$, where in the run for $\mathcal{J}'$ the new job $j^*$ is inserted to the list in any position consistent with LPT. This is without loss of generality since different tie breaking do not affect the load profiles of the solutions.

Consider the set of instances $(\mathcal{J}|_k, \mathcal{M})$ for $k = r, \ldots, n$, where $\mathcal{J}|_k \subseteq \mathcal{J}$ is the set of the $k$ largest jobs in $\mathcal{J}$, and $r$ is the maximal index such that $\tilde{p}_r \leq \tilde{p}_{j^*}$. Similarly, let $\mathcal{J}'|_k = \mathcal{J}|_k \cup \{j^*\}$ for any $k \in \{r, \ldots, n\}$. We will show by induction that the lemma is true for each pair $(\mathcal{J}|_k, \mathcal{M})$ and $(\mathcal{J}'|_k, \mathcal{M})$. The base case $k = r$ follows easily from Lemma 6 since $\mathcal{S}_{\text{LPT}}(\mathcal{J}|_k, \mathcal{M})$ and $\mathcal{S}_{\text{LPT}}(\mathcal{J}'|_k \setminus \{j^*\}, \mathcal{M})$ assign to the same machines all jobs $\{1, \ldots, r\}$, and adding $j^*$ to the least loaded machine in $\mathcal{S}_{\text{LPT}}(\mathcal{J}'|_k \setminus \{j^*\}, \mathcal{M})$ (and a job of size 0 to the least loaded machine in $\mathcal{S}_{\text{LPT}}(\mathcal{J}|_k, \mathcal{M})$) is the same as adding $\tilde{p}_{j^*}$ to the first coordinate of $\text{load}(\mathcal{S}_{\text{LPT}}(\mathcal{J}'|_k \setminus \{j^*\}, \mathcal{M}))$, and then the inequality holds.

Suppose now that $\text{load}(\mathcal{S}_{\text{LPT}}(\mathcal{J}|_k, \mathcal{M})) \leq \text{load}(\mathcal{S}_{\text{LPT}}(\mathcal{J}'|_k, \mathcal{M}))$. Showing that the inequality is true for $k+1$ is equivalent to show that when assigning job $k+1$ to a least loaded machine in $\mathcal{S}_{\text{LPT}}(\mathcal{J}|_k, \mathcal{M})$ and in $\mathcal{S}_{\text{LPT}}(\mathcal{J}'|_k, \mathcal{M})$, the resulting load profiles satisfy the inequality, which is precisely the statement of Lemma 6 adding $\tilde{p}_{k+1}$ to the first coordinate of $\text{load}(\mathcal{S}_{\text{LPT}}(\mathcal{J}|_k, \mathcal{M}))$ and also to the first coordinate of $\text{load}(\mathcal{S}_{\text{LPT}}(\mathcal{J}'|_k, \mathcal{M}))$. □

This lemma together with our rounding procedure allow us to show that the difference (in terms of the Hamming distance) of the load profiles of two consecutive solutions consisting purely of big jobs, is bounded by a small constant.

**Lemma 8.** *Consider two instances $(\mathcal{J}, \mathcal{M})$ and $(\mathcal{J}', \mathcal{M})$ with $\mathcal{J}' = \mathcal{J} \cup \{j^*\}$, where $\mathcal{J}'$ contains only big or huge jobs w.r.t* UB. *Then the vectors* $\text{load}(\mathcal{S}_{\text{LPT}}(\mathcal{J}, \mathcal{M}))$ *and* $\text{load}(\mathcal{S}_{\text{LPT}}(\mathcal{J}', \mathcal{M}))$ *differ in at most $\frac{\tilde{p}_{j^*}}{\varepsilon 2^\ell} \in O(1/\varepsilon^2)$ many coordinates.*

*Proof.* Due to Lemma 7 we have that

$$\text{load}(\mathcal{S}_{\text{LPT}}(\mathcal{J}, \mathcal{M})) = (t_1, \ldots, t_m) \leq (t'_1, \ldots, t'_m) = \text{load}(\mathcal{S}_{\text{LPT}}(\mathcal{J}', \mathcal{M})).$$

Also, if $t_i < t'_i$ for some $i$, then $t'_i \geq t_i + \varepsilon 2^\ell$ since all values $t_j, t_{j'}$ are integer multiples of $\varepsilon 2^\ell$ because of Lemma 4. Since $||\text{load}(\mathcal{S}_{\text{LPT}}(\mathcal{J}', \mathcal{M})) - \text{load}(\mathcal{S}_{\text{LPT}}(\mathcal{J}, \mathcal{M}))||_1 = \tilde{p}_{j^*}$, we obtain that the number of coordinates in which the load profiles differ is at most $\frac{\tilde{p}_{j^*}}{\varepsilon 2^\ell}$. Finally, recalling that $j^*$ is big, then $\tilde{p}_{j^*} \leq 2^u \leq \text{UB} \leq 2^\ell/\varepsilon$, and we can bound the number of different coordinates by $\frac{\tilde{p}_{j^*}}{\varepsilon 2^\ell} \leq 1/\varepsilon^2 \in O(1/\varepsilon^2)$. □

**Description of Online LPT.** Consider two instances $(\mathcal{J}, \mathcal{M})$ and $(\mathcal{J}', \mathcal{M})$ such that $\mathcal{J}' = \mathcal{J} \cup \{j^*\}$, and let OPT and OPT$'$ be their optimal values, respectively. We say that an assignment is an LPT-solution if there is some tie-breaking rule such that LPT yields such assignment. We will compute an upper bound UB on OPT$'$ by computing an LPT-solution and duplicating the value of its minimum load. For this upper bound we compute its respective set $\tilde{P}$ with (1) and (2). In the algorithm we will label elements in $\tilde{P} = \{q_1, \ldots, q_{|\tilde{P}|}\}$ such that $q_1 > q_2 > \cdots > q_{|\tilde{P}|}$. Let $\mathcal{J}_h \subseteq \mathcal{J}$ (respectively $\mathcal{J}'_h \subseteq \mathcal{J}'$) be the set of jobs of size $q_h$ in $\mathcal{J}$ (respectively $\mathcal{J}'$), for $q_h \in \tilde{P}$. Similarly, we define $\mathcal{J}_0$ (resp. $\mathcal{J}'_0$) to be the set of jobs in $\mathcal{J}$ (resp. $\mathcal{J}'$) of sizes larger than $q_1$, that is, all huge jobs in $\mathcal{J}$ (resp. $\mathcal{J}'$). Also, let $\mathcal{S}_h$ (resp. $\mathcal{S}'_h$) be the solution $\mathcal{S}$ (resp. $\mathcal{S}'$) restricted to jobs of size $q_h$ or larger. Finally, $\mathcal{S}_0$ and $\mathcal{S}'_0$ are the respective solutions restricted to jobs in $\mathcal{J}_0$.

In the description of the algorithm we use the concept of a *list-scheduling* algorithm, that refers to assigning jobs iteratively (in any order) to some machine of minimum load. In what follows, $x_+$ denotes the positive part of $x \in \mathbb{R}$, i.e., $x_+ = \max\{x, 0\}$.

To understand the algorithm, it is useful to have the following observation in mind.

**Observation 9.** *Consider a solution $\mathcal{S}$ for jobs in $\mathcal{J}$ and let $\mathcal{K}$ be a set of jobs with $\mathcal{J} \cap \mathcal{K} = \emptyset$ and all jobs in $\mathcal{K}$ have the same size $p$. Consider a solution $\mathcal{S}_{LS}$ constructed by adding the jobs from $\mathcal{K}$ in $\mathcal{S}$ using list-scheduling, and let $\lambda = \ell_{\min}(\mathcal{S}_{LS})$. Notice that $\lambda$ is independent of the tie-breaking rule used in list-scheduling. Consider any solution $\mathcal{S}'$ that is constructed starting from $\mathcal{S}$ and adding jobs in $\mathcal{K}$ in some arbitrary way. Then $\mathcal{S}'$ corresponds to a solution obtained by adding jobs from $\mathcal{K}$ with a list-scheduling procedure (for some tie-breaking rule) if and only if the number of jobs in $\mathcal{K}$ added to each machine $i$ is: (i) $\left\lceil \frac{(\lambda - \ell_i(\mathcal{S}))_+}{p} \right\rceil$ if $\frac{(\lambda - \ell_i(\mathcal{S}))_+}{p}$ is not an integer, and either $\frac{(\lambda - \ell_i(\mathcal{S}))_+}{p}$ or $\frac{(\lambda - \ell_i(\mathcal{S}))_+}{p} + 1$ if $\frac{(\lambda - \ell_i(\mathcal{S}))_+}{p}$ is a non-negative integer.*

Our main procedure is called every time that we get a new job $j^*$ (where $\mathcal{J}' = \mathcal{J} \cup \{j^*\}$) and receives as input the current solution $\mathcal{S}$ for $(\mathcal{J}, \mathcal{M})$. If $\mathcal{J} = \emptyset$ then $\mathcal{S}$ is trivially initialized as empty. Broadly speaking, it works in phases $h \in \{0, \ldots, |\tilde{P}|\}$, where for each $h$ it assigns jobs in $\mathcal{J}'_h$. First, we assign jobs exactly as in $\mathcal{S}_h$ for machines in which the assignment of $\mathcal{S}_{h-1}$ and $\mathcal{S}'_{h-1}$ coincide. As we will see, this is consistent with LPT by the previous observation and Lemma 7. The remaining jobs in $\mathcal{J}'_h$ are assigned using list-scheduling. Crucially, we will break ties in favor of machines where the assignment of $\mathcal{S}_{h-1}$ and $\mathcal{S}'_{h-1}$ differ. This is necessary to avoid creating new machines with different assignments. After assigning huge and big jobs, small jobs are added exactly as in $\mathcal{S}$ in machines where the assignment of big jobs in $\mathcal{S}$ and $\mathcal{S}'$ coincides. The rest of small jobs are added greedily. In the last part, the algorithm rebalances small jobs by moving them from machines of load higher than $\ell_i(\mathcal{S}') + 2^\ell$ to the least loaded machines. The exact description is given in Algorithm 1.

During the analysis we will consider the run of the algorithm for solution $\mathcal{S}$. For such run we will call $\text{UB}_0$ the corresponding upper bound and $\tilde{P}_0$ the set of big job sizes as defined by (1) and (2). Similarly we will denote by $\ell_0$ and $u_0$ the smallest and largest values in $I(\text{UB}_0)$, as defined in (1). First we prove the following auxiliary lemma.

**Lemma 10.** *If $\mathcal{S}'$ is the output of the algorithm then $\mathcal{S}'_{|\tilde{P}|}$ is an LPT-solution.*

*Proof.* We show the proof inductively. Consider a run of the algorithm with input assignment $\mathcal{S}$. If $\mathcal{S}$ is empty then it is clearly an LPT-solution. Otherwise, $\mathcal{S}$ is the output of a run of the algorithm. We can assume inductively that $\mathcal{S}_{|\tilde{P}_0|}$ is an LPT-solution (and thus also any restriction of $\mathcal{S}_{|\tilde{P}_0|}$ to

**Algorithm 1** Online LPT

**Input:** Instances $(\mathcal{J}, \mathcal{M})$ and $(\mathcal{J}', \mathcal{M})$ such that $\mathcal{J}' = \mathcal{J} \cup \{j^*\}$; a schedule $\mathcal{S}(\mathcal{J}, \mathcal{M})$.

1: run LPT on input $\mathcal{J}'$ and let $\tau$ be the minimum load of the constructed solution. Set $\text{UB} \leftarrow 2\tau$. Define $\tilde{P}, \ell$, and $u$ based on this upper bound UB using (1) and (2).
2: set $\mathcal{M}^{=}_{-1} \leftarrow \mathcal{M}$ and $\mathcal{M}^{\neq}_{-1} \leftarrow \emptyset$.
3: **for** $h = 0, 1, \ldots, |\tilde{P}|$ **do** ▷ `Assignment of big and huge jobs`
4: for each machine $i \in \mathcal{M}^{=}_{h-1}$, assign all jobs in $\mathcal{J}_h \cap \mathcal{S}^{-1}(i)$ to $i$ in $\mathcal{S}'$.
5: for jobs in $\mathcal{J}'_h$ still not assigned in $\mathcal{S}'$, apply list-scheduling (with an arbitrary order of jobs). If there is more than one least loaded machine break ties in favor of machines in $\mathcal{M}^{\neq}_{h-1}$.
6: define $\mathcal{M}^{=}_h$ as the set of machines $i$ such that $\mathcal{S}_h^{-1}(i) = \mathcal{S}_h'^{-1}(i)$ and $\mathcal{M}^{\neq}_h \leftarrow \mathcal{M} \setminus \mathcal{M}^{=}_h$.
7: **end for**
8: **for** machines $i \in \mathcal{M}^{=}_{|\tilde{P}|}$ **do** ▷ `Assignment of small jobs`
9: assign all small jobs w.r.t to UB in $\mathcal{J} \cap \mathcal{S}^{-1}(i)$ to $i$ in $\mathcal{S}'$.
10: **end for**
11: assign the remaining jobs using list-scheduling.
12: set $\overline{\mathcal{M}}$ to be the set of machines containing a small job w.r.t UB.
13: **while** there exists $i \in \overline{\mathcal{M}}$ s.t. $\ell_i(\mathcal{S}') > \ell_{\min}(\mathcal{S}') + 2^{\ell}$ **do**
14: consider a machine $i \in \overline{\mathcal{M}}$ of maximum load. Reassign the smallest job in $\mathcal{S}'^{-1}(i)$ to any least loaded machine.
15: update $\overline{\mathcal{M}}$ to be the set of machines containing a small job w.r.t UB.
16: **end while**
17: **return** $\mathcal{S}'$.

jobs of sizes at least $p$, for any $p \geq 0$). Notice that $\text{UB}_0 \leq \text{UB}$, by Lemma 7, and thus $\min \tilde{P}_0 \leq \min \tilde{P}$ and $\max \tilde{P}_0 \leq \max \tilde{P}$.

We use a second induction to show that, for every $h \in \{0, \ldots, |\tilde{P}|\}$, $\mathcal{S}'_h$ is an LPT-solution .

To show the base case ($h = 0$), consider jobs in $J'_0$, which are all larger than $\text{UB} \geq \text{OPT}'$. Hence there are at most $m$ of them, and the algorithm assigns them each to a different machine (this, again, follows inductively). Thus, the base case holds.

Consider $h \geq 1$ and let us assume that $\mathcal{S}'_{h-1}$ is an LPT-solution. Let $\mathcal{S}_{\text{LPT},h}$ be an LPT-solution for jobs in $\mathcal{J}_0 \cup \ldots \cup \mathcal{J}_h$, and similarly $\mathcal{S}'_{\text{LPT},h}$ for jobs in $\mathcal{J}'_0 \cup \ldots \cup \mathcal{J}'_h$. First observe that the load profile vector $\text{load}(\mathcal{S}'_{\text{LPT},h})$ is independent of the tie-breaking rule. Consider the target value $\lambda = \ell_{\min}(\mathcal{S}_{\text{LPT},h})$ and $\lambda' = \ell_{\min}(\mathcal{S}'_{\text{LPT},h})$. Notice that by Lemma 7 $\lambda \leq \lambda'$.

Since $\mathcal{S}_{h-1}$ is an LPT-solution, then $\mathcal{S}'_h$ is an LPT-solution if jobs in $\mathcal{J}'_h$ are added using list-scheduling. By Observation 9 the following characterizes this fact: for all machines $i \in M$, the number of jobs assigned in $\mathcal{S}'_h$ to $i$ is: (i) $\lceil (\lambda' - \ell_i(\mathcal{S}'_{h-1}))_+ / q_h \rceil$ if $(\lambda' - \ell_i(\mathcal{S}'_{h-1}))_+ / q_h$ is not an integer, and either $(\lambda' - \ell_i(\mathcal{S}'_{h-1}))_+ / q_h$ or $(\lambda' - \ell_i(\mathcal{S}'_{h-1}))_+ / q_h + 1$ if $(\lambda' - \ell_i(\mathcal{S}'_{h-1}))_+ / q_h$ is an integer. Since $\lambda \leq \lambda'$, and $\mathcal{S}_h$ is an LPT-solution, then the number of jobs assigned in Step 4 is never more than $\lceil (\lambda' - \ell_i(\mathcal{S}'_{h-1}))_+ / q_h \rceil$ if $(\lambda' - \ell_i(\mathcal{S}'_{h-1}))_+ / q_h$ is not an integer, and never more than $(\lambda' - \ell_i(\mathcal{S}'_{h-1}))_+ / q_h + 1$ if $(\lambda' - \ell_i(\mathcal{S}'_{h-1}))_+ / q_h$ is integer. This implies that after adding jobs in Step 5 we obtain an LPT-solution. □

Now we can argue about the approximation guarantee of the obtained solution.

**Lemma 11.** *The algorithm is* $(4/3 + O(\varepsilon))$*-competitive.*

*Proof.* We will use the previous lemma to show that $\mathcal{S}'$ is a $k$-relaxed version of $\mathcal{S}_{\text{LPT}}(\mathcal{J}', \mathcal{M})$ for some $k \leq 4$, which is enough to conclude the lemma due to Lemma 3. Indeed, let $k = 2^\ell/(\varepsilon \text{OPT}')$. Then, by the previous lemma all jobs larger than $k\varepsilon\text{OPT}' = 2^\ell$ are assigned with LPT. Also, the while loop at Step 13 ensures that the output of the algorithm is a $k$-relaxed version of $\mathcal{S}_{\text{LPT}}(\mathcal{J}', \mathcal{M})$. The lemma follows since $k = 2^\ell/(\varepsilon\text{OPT}') \leq 2\varepsilon\text{UB}/(\varepsilon\text{OPT}') \leq 4$. □

**Bounding the migration factor.** To analyze the migration factor of the algorithm, we will show that $|\mathcal{M}^{\neq}_{|\tilde{P}|}|$ is upper bounded by a constant. This will be done inductively by first bounding $|\mathcal{M}^{\neq}_h \setminus \mathcal{M}^{\neq}_{h-1}|$ for each $h$ and then using the fact that $|\tilde{P}| \in O((1/\varepsilon)\log(1/\varepsilon))$. To understand the overall idea of the proof consider Figure 2.

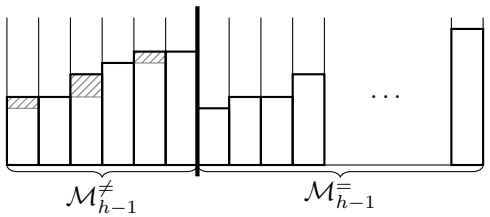


Figure 2: Depiction of a possible situation at the end of iteration $h-1$. The machines on the right side correspond to machines in $\mathcal{M}^{=}_{h-1}$ and therefore process the same jobs in $\mathcal{S}_{h-1}$ and $\mathcal{S}'_{h-1}$. Assume, possibly erroneously and just as a thought experiment, that the machines in $\mathcal{M}^{\neq}_{h-1}$ can be sorted non-decreasingly by load for $\mathcal{S}_{h-1}$ and $\mathcal{S}'_{h-1}$ simultaneously. The two solutions are depicted simultaneously in the picture, where the difference of loads on machines in $\mathcal{M}^{\neq}_{h-1}$ corresponds to the dashed area. This load equals to $\tilde{p}_{j^*}$, which is spread in only constantly many machines by Lemma 8. When assigning jobs in $\mathcal{J}_h$, the algorithm first assigns a number of jobs to each machine in $\mathcal{M}^{=}_{h-1}$ (Step 4), and then fills machines in $\mathcal{M}^{\neq}_{h-1}$. Notice that while the algorithm does not assign another job to a machine in $\mathcal{M}^{=}_{h-1}$, no new machine will enter $\mathcal{M}^{\neq}_h \setminus \mathcal{M}^{\neq}_{h-1}$. On the other hand, the number of such jobs can be bounded by a number proportional to $\tilde{p}_{j^*}$ (and $1/\varepsilon$), which then also bounds the number of machines in $\mathcal{M}^{\neq}_h \setminus \mathcal{M}^{\neq}_{h-1}$. In reality, however, it is not true that the machines in $\mathcal{M}^{\neq}_{h-1}$ can be sorted non-decreasingly on the loads for $\mathcal{S}_{h-1}$ and $\mathcal{S}'_{h-1}$ simultaneously. This provokes a number of technical difficulties that we avoid by using a different permutation of machines for each solution and invoking Lemma 7.

Let us consider huge jobs w.r.t UB (i.e. jobs in $\mathcal{J}'_0$). Notice that all these jobs are larger than $\text{OPT}' \geq \text{OPT}$, and hence in $\mathcal{S}'_0$ each one is assigned alone to one machine. The same situation happens in solution $\mathcal{S}$ restricted to jobs in $\mathcal{J}_0$. Thus, none of these jobs are migrated. Hence, we can assume w.l.o.g. for the sake of the analysis of the migration that all jobs are big or small w.r.t UB (including $j^*$). Additionally, we can assume that $j^*$ is not small, since otherwise there is no migration.

As in the proof of Lemma 10, let $\mathcal{S}_{\text{LPT},h}$ be an LPT-solution for jobs in $\mathcal{J}_0 \cup \ldots \cup \mathcal{J}_h$, and similarly $\mathcal{S}'_{\text{LPT},h}$ for jobs in $\mathcal{J}'_0 \cup \ldots \cup \mathcal{J}'_h$. Let us fix $h \geq 1$ and consider the target values $\lambda = \ell_{\min}(\mathcal{S}_{\text{LPT},h})$ and $\lambda' = \ell_{\min}(\mathcal{S}'_{\text{LPT},h})$. Notice that by Lemma 7, $\lambda \leq \lambda'$. Let $\mathcal{J}^{=}_h$ be the set of jobs assigned by Step 5 to machines in $\mathcal{M}^{=}_{h-1}$. Notice that the jobs in $\mathcal{J}^{=}_h$ correspond to the jobs in $\mathcal{J}'_h$ that $\mathcal{S}'$ assigns to a machine in $\mathcal{M}^{=}_{h-1}$ but $\mathcal{S}$ processes in $\mathcal{M}^{\neq}_{h-1}$.

Our strategy will be to bound the cardinality of set $\mathcal{J}^{=}_h$ and then use this to upper bound $|\mathcal{M}^{\neq}_h \setminus \mathcal{M}^{\neq}_{h-1}|$. First we prove two auxiliary lemmas that help to upper bound $|\mathcal{J}^{=}_h|$.

**Lemma 12.** *Assume that $\mathcal{J}^{=}_h \neq \emptyset$. For each machine $i \in \mathcal{M}^{\neq}_{h-1}$, if $\lambda - \ell_i(\mathcal{S}'_{h-1}) \geq 0$ solution $\mathcal{S}'_h$ assigns to $i$ at least $\left\lfloor \frac{(\lambda - \ell_i(\mathcal{S}'_{h-1}))_+}{q_h} \right\rfloor + 1$ many jobs from $\mathcal{J}_h$.*

*Proof.* We consider two cases. If $\lambda' = \lambda$, then the number of jobs in $\mathcal{J}_h$ assigned to machine $i$ is at least $\left\lfloor \frac{(\lambda' - \ell_i(\mathcal{S}'_{h-1}))_+}{q_h} \right\rfloor + 1$. Indeed, if $\frac{(\lambda' - \ell_i(\mathcal{S}'_{h-1}))_+}{q_h}$ is fractional then the number of jobs must be $\left\lceil \frac{(\lambda' - \ell_i(\mathcal{S}'_{h-1}))_+}{q_h} \right\rceil = \left\lfloor \frac{(\lambda' - \ell_i(\mathcal{S}'_{h-1}))_+}{q_h} \right\rfloor + 1$. If, on the other hand, $\frac{(\lambda' - \ell_i(\mathcal{S}'_{h-1}))_+}{q_h}$ is integral, then the algorithm might assign to $i \in \mathcal{M}^{\neq}_{h-1}$ only $\frac{(\lambda' - \ell_i(\mathcal{S}'_{h-1}))_+}{q_h} = \left\lfloor \frac{(\lambda - \ell_i(\mathcal{S}'_{h-1}))_+}{q_h} \right\rfloor$ many jobs. However, if this is the case the tie-breaking rule in Step 5 implies that $\mathcal{J}^{=}_h = \emptyset$, which contradicts our hypothesis. Then the number of assigned jobs is exactly $\left\lfloor \frac{(\lambda' - \ell_i(\mathcal{S}'_{h-1}))_+}{q_h} \right\rfloor + 1$, and thus the claim holds.

If $\lambda' > \lambda$, then $\lambda' > \ell_i(\mathcal{S}'_{h-1})$ and the number of jobs in $\mathcal{J}_h$ assigned to machine $i$ is at least $\left\lceil \frac{(\lambda' - \ell_i(\mathcal{S}'_{h-1}))_+}{q_h} \right\rceil = \left\lceil \frac{\lambda' - \ell_i(\mathcal{S}'_{h-1})}{q_h} \right\rceil \geq \left\lfloor \frac{(\lambda - \ell_i(\mathcal{S}'_{h-1}))}{q_h} \right\rfloor + 1 = \left\lfloor \frac{(\lambda - \ell_i(\mathcal{S}'_{h-1}))_+}{q_h} \right\rfloor + 1$ and hence the claim holds. □

**Lemma 13.** *Let $x, y \in \mathbb{R}^n_+$, $x = (x_1, x_2, \ldots, x_n)$, $y = (y_1, y_2, \ldots, y_n)$ such that $x_1 \leq x_2 \leq \cdots \leq x_n$, $y_1 \leq y_2 \leq \cdots \leq y_n$ and $x \leq y$ coordinate-wise. Assume that $x_j = y_i$ for some indices $i, j$. If $x_{-i}$ denotes the $(n-1)$-dimensional vector obtained by removing the $i$-th entry of $x$, and $y_{-j}$ is the vector obtained by removing the $j$-th entry of $y$, then $x_{-i} \leq y_{-j}$.*

*Proof.* Notice first that if $i = j$ then the result is a direct consequence of Lemma 6: by taking $\alpha = \beta = -x_i$ and coordinate $i$, we get new vectors $\tilde{x}$ and $\tilde{y}$ satisfying $\tilde{x} \leq \tilde{y}$ and $\tilde{x}_1 = \tilde{y}_1 = 0$, and hence we can conclude that $x_{-i} \leq y_{-j}$ because $x_{-i}$ (resp $y_{-j}$) corresponds to the last $n-1$ coordinates of $\tilde{x}$ (resp. $\tilde{y}$).

We now distinguish two cases: if $i < j$, we have that $y_j = x_i \leq x_j \leq y_j$, hence $x_k = y_j$ for every $k = i, i+1, \ldots, j$. This implies that $x_{-i} = x_{-(i+1)} = \cdots = x_{-j}$, and then we can conclude that $x_{-i} \leq y_{-j}$ by applying the previous observation for $x_{-j}$ and $y_{-j}$. On the other hand, if $j < i$, we define vector $z$ equal to $y$ but replacing coordinates $j, j+1, \ldots, i$ by $y_j$. It is not difficult to see that $x \leq z \leq y$ coordinate-wise, and also $z_{-j} = z_{-(j+1)} = \cdots = z_{-i}$. If we apply the first observation for $x$ and $z$ using coordinate $i$ we have that $x_{-i} \leq z_{-i}$, and applying it to $z$ and $y$ using coordinate $j$ we get that $z_{-j} \leq y_{-j}$. Merging both inequalities and using the fact that $z_{-i} = z_{-j}$, we conclude that $x_{-i} \leq y_{-j}$. □

**Lemma 14.** *It holds that $|\mathcal{J}^{=}_h| \in O(\frac{\tilde{p}_{j^*}}{\varepsilon 2^{\ell}})$.*

*Proof.* Assume, w.l.o.g., that $\mathcal{M}^{\neq}_{h-1} = \{1, \ldots, m'\}$ and that $\ell_1(\mathcal{S}'_{h-1}) \leq \ell_2(\mathcal{S}'_{h-1}) \leq \ldots \leq \ell_{m'}(\mathcal{S}'_{h-1})$. Consider also a permutation $\sigma : \mathcal{M}^{\neq}_{h-1} \to \mathcal{M}^{\neq}_{h-1}$ such that $\ell_{\sigma(1)}(\mathcal{S}_{h-1}) \leq \ell_{\sigma(2)}(\mathcal{S}_{h-1}) \leq \ldots \leq \ell_{\sigma(m')}(\mathcal{S}_{h-1})$. By Lemma 7 the sorted vector of loads (over all machines) of solution $\mathcal{S}_{h-1}$ is upper bounded by the sorted vector of loads of $\mathcal{S}'_{h-1}$. Applying Lemma 13 iteratively to remove machines in $\mathcal{M}^{=}_{h-1}$ one by one (which have the same assignment in both solutions), it holds that $\ell_{\sigma(i)}(\mathcal{S}_{h-1}) \leq \ell_i(\mathcal{S}'_{h-1})$ for all $i \in \mathcal{M}^{\neq}_{h-1}$. Let us consider sets

$$T_- = \{i \in \mathcal{M}^{\neq}_{h-1} : \ell_i(\mathcal{S}'_{h-1}) \leq \lambda\}, \text{ and}$$
$$T_+ = \{i \in \mathcal{M}^{\neq}_{h-1} : \ell_{\sigma(i)}(\mathcal{S}_{h-1}) \leq \lambda \text{ and } \ell_i(\mathcal{S}'_{h-1}) > \lambda\}.$$

Lemma 12 implies that the total number of jobs from $\mathcal{J}'_h$ assigned by $\mathcal{S}'_h$ to machines in $\mathcal{M}^{\neq}_{h-1}$

is at least

$$\begin{aligned}
&\sum_{i\in T_-}\left(\left\lfloor\tfrac{(\lambda-\ell_i(\mathcal{S}'_{h-1}))_+}{q_h}\right\rfloor+1\right)\\
&=\sum_{i\in T_-}\left(\left\lfloor\tfrac{(\lambda-\ell_{\sigma(i)}(\mathcal{S}_{h-1}))_+}{q_h}\right\rfloor+1\right)+\sum_{i\in T_-}\left\lfloor\tfrac{(\lambda-\ell_i(\mathcal{S}'_{h-1}))_+}{q_h}\right\rfloor-\left\lfloor\tfrac{(\lambda-\ell_{\sigma(i)}(\mathcal{S}_{h-1}))_+}{q_h}\right\rfloor\\
&=\sum_{i\in T_-\cup T_+}\left(\left\lfloor\tfrac{(\lambda-\ell_{\sigma(i)}(\mathcal{S}_{h-1}))_+}{q_h}\right\rfloor+1\right)-\sum_{i\in T_+}\left(\left\lfloor\tfrac{(\lambda-\ell_{\sigma(i)}(\mathcal{S}_{h-1}))_+}{q_h}\right\rfloor+1\right)\\
&\quad+\sum_{i\in T_-}\left\lfloor\tfrac{(\lambda-\ell_i(\mathcal{S}'_{h-1}))_+}{q_h}\right\rfloor-\left\lfloor\tfrac{(\lambda-\ell_{\sigma(i)}(\mathcal{S}_{h-1}))_+}{q_h}\right\rfloor.
\end{aligned}$$

Notice that the set $T_-\cup T_+$ contains all indices $i\in\mathcal{M}^{\neq}_{h-1}$ such that $\ell_{\sigma(i)}(\mathcal{S}_{h-1})\leq\lambda$. Hence, the first sum in the last expression upper bounds the number of jobs in $\mathcal{J}_h$ that solution $\mathcal{S}$ assigns to machines in $\mathcal{M}^{\neq}_{h-1}$. Hence, since $|\mathcal{J}'_h\setminus\mathcal{J}_h|\leq 1$, it holds that

$$\begin{aligned}
|\mathcal{J}^{=}_h|&\leq 1+\sum_{i\in T_+}\left(\left\lfloor\tfrac{(\lambda-\ell_{\sigma(i)}(\mathcal{S}_{h-1}))_+}{q_h}\right\rfloor+1\right)+\sum_{i\in T_-}\left\lfloor\tfrac{(\lambda-\ell_{\sigma(i)}(\mathcal{S}_{h-1}))_+}{q_h}\right\rfloor-\left\lfloor\tfrac{(\lambda-\ell_i(\mathcal{S}'_{h-1}))_+}{q_h}\right\rfloor\\
&\leq 1+|T_+|+\sum_{i\in T_+}\left\lfloor\tfrac{(\lambda-\ell_{\sigma(i)}(\mathcal{S}_{h-1}))_+}{q_h}\right\rfloor+\sum_{i\in T_-}\left\lfloor\tfrac{(\lambda-\ell_{\sigma(i)}(\mathcal{S}_{h-1}))_+}{q_h}\right\rfloor-\left\lfloor\tfrac{(\lambda-\ell_i(\mathcal{S}'_{h-1}))_+}{q_h}\right\rfloor\\
&\leq 1+|T_+|+\underbrace{\sum_{i\in T_+}\left\lfloor\tfrac{(\lambda-\ell_i(\mathcal{S}'_{h-1}))_+}{q_h}\right\rfloor}_{=0}+\sum_{i\in T_-\cup T_+}\left\lfloor\tfrac{(\lambda-\ell_{\sigma(i)}(\mathcal{S}_{h-1}))_+}{q_h}\right\rfloor-\left\lfloor\tfrac{(\lambda-\ell_i(\mathcal{S}'_{h-1}))_+}{q_h}\right\rfloor.
\end{aligned}$$

Let us now consider $T_{\neq}=\{i\in\mathcal{M}^{\neq}_{h-1}:\ell_{\sigma(i)}(\mathcal{S}_{h-1})\neq\ell_i(\mathcal{S}'_{h-1})\}$. Thus, the last expression is at most

$$\begin{aligned}
|\mathcal{J}^{=}_h|&\leq 1+|T_+|+\sum_{i\in(T_-\cup T_+)\cap T_{\neq}}\left\lfloor\tfrac{(\lambda-\ell_{\sigma(i)}(\mathcal{S}_{h-1}))_+}{q_h}\right\rfloor-\left\lfloor\tfrac{(\lambda-\ell_i(\mathcal{S}'_{h-1}))_+}{q_h}\right\rfloor\\
&\leq 1+|T_+|+\sum_{i\in T_{\neq}}\left(\tfrac{(\lambda-\ell_{\sigma(i)}(\mathcal{S}_{h-1}))_+}{q_h}-\tfrac{(\lambda-\ell_i(\mathcal{S}'_{h-1}))_+}{q_h}+1\right)\\
&\leq 1+|T_+|+|T_{\neq}|+\sum_{i\in T_{\neq}}\tfrac{\ell_i(\mathcal{S}'_{h-1})-\ell_{\sigma(i)}(\mathcal{S}_{h-1})}{q_h}\\
&\leq 1+2|T_{\neq}|+\tfrac{\tilde{p}_{j^*}}{q_h}.
\end{aligned}$$

Also, Lemma 8 can be applied and thus $|T_{\neq}|\leq\frac{\tilde{p}_{j^*}}{\varepsilon 2^{\ell}}$. The lemma finally follows since $q_h\geq 2^{\ell}$ by definition. □

Notice that jobs in $\mathcal{J}^{=}_h$ are the only jobs assigned in a given iteration $h$ that can cause one new machine to have different assignments in $\mathcal{S}_h$ and $\mathcal{S}'_h$. Thus, $|\mathcal{M}^{\neq}_h\setminus\mathcal{M}^{\neq}_{h-1}|\leq|\mathcal{J}^{=}_h|$ and the following lemma holds.

**Lemma 15.** *For all $h \in \{1, \dots, |\tilde{P}|\}$ it holds that $|\mathcal{M}^{\neq}_h \setminus \mathcal{M}^{\neq}_{h-1}| \in O(\frac{\tilde{p}_{j^*}}{\varepsilon 2^\ell})$.*

Putting all the discussed ideas together, we prove the following result.

**Theorem 16.** *Algorithm Online LPT is a polynomial time $(4/3+O(\varepsilon))$-competitive algorithm with $O((1/\varepsilon^3)\log(1/\varepsilon))$ migration factor.*

*Proof.* We first argue that the algorithm runs in polynomial time. Indeed, it suffices to show that the algorithm enters the while loop in Step 13 a polynomial number of times. This follows easily as the quantity $\ell_{\min}(\mathcal{S}')$ is non-decreasing, and hence a job can be reassigned to a least loaded machine at most once .

Let us now bound the migration factor. We do this in two steps. First consider solution $\mathcal{S}'$ before entering Step 13. We first bound the volume of jobs migrated between $\mathcal{S}$ and $\mathcal{S}'$, and then bound the total volume of jobs reassigned in the while loop in Step 13.

For the first bound, by the previous lemma and since $\mathcal{M}^{\neq}_{-1} = \emptyset$, it holds that $|\mathcal{M}^{\neq}_{|\tilde{P}|}| \leq |\tilde{P}| \cdot O(\frac{\tilde{p}_{j^*}}{\varepsilon 2^\ell}) \leq O((1/\varepsilon)\log(1/\varepsilon)\frac{\tilde{p}_{j^*}}{\varepsilon 2^\ell})$. The load of jobs in $\mathcal{S}'_{|\tilde{P}|}$ that are migrated is upper bounded by $\sum_{i \in \mathcal{M}^{\neq}_{|\tilde{P}|}} \ell_i(\mathcal{S}'_{|\tilde{P}|}) \leq |\mathcal{M}^{\neq}_{|\tilde{P}|}| \max_{i \in \mathcal{M}^{\neq}_{|\tilde{P}|}} \ell_i(\mathcal{S}'_{|\tilde{P}|})$. On the other hand, since we are assuming (w.l.o.g) that there is no huge job, the total load of each machine is at most 2UB as argued in Section 2.2. We conclude that the big jobs migrated have a total load of at most $2\text{UB} \cdot |\mathcal{M}^{\neq}_{|\tilde{P}|}| = \text{UB} \cdot O((1/\varepsilon)\log(1/\varepsilon)\frac{\tilde{p}_{j^*}}{\varepsilon 2^\ell})$. Finally, notice that small jobs migrated (before entering Step 13) are the ones assigned to machines in $\mathcal{M}^{\neq}_{|\tilde{P}|}$ by $\mathcal{S}$. Since $\mathcal{S}$ is the output of Online LPT, then the total load of these jobs is at most $(\ell_{\min}(\mathcal{S}') + 2^\ell) \cdot \mathcal{M}^{\neq}_{|\tilde{P}|} \leq 2\text{UB} \cdot \mathcal{M}^{\neq}_{|\tilde{P}|} \leq \text{UB} \cdot O((1/\varepsilon)\log(1/\varepsilon)\frac{\tilde{p}_{j^*}}{\varepsilon 2^\ell})$. We conclude that the total load migrated is at most $\text{UB} \cdot O((1/\varepsilon)\log(1/\varepsilon)\frac{\tilde{p}_{j^*}}{\varepsilon 2^\ell})$.

It remains to bound the volume migrated in the while loop of Step 13. For this we will show the following claim.

**Claim:** Let $\mathcal{S}'$ be the solution constructed before entering Step 13. Then all reassigned jobs in the while loop, except possibly the one reassigned last, are assigned to a machine in $\mathcal{M}^{\neq}_{|\tilde{P}|}$ by $\mathcal{S}'$.

Assume the claim holds and let us consider the solution $\mathcal{S}'$ as output by the algorithm. Then the total volume of reassigned jobs is bounded by $|\mathcal{M}^{\neq}_{|\tilde{P}|}| \max_{i \in \mathcal{M}^{\neq}_{|\tilde{P}|}} \ell_i(\mathcal{S}')$. Since by construction the load of a machine that process a job smaller than $2^\ell$ is at most $\ell_{\min}(\mathcal{S}') + 2^\ell \leq 2\text{UB}$, the total volume migrated will be at most $\text{UB} \cdot O((1/\varepsilon)\log(1/\varepsilon)\frac{\tilde{p}_{j^*}}{\varepsilon 2^\ell})$ as before. Hence, the migration factor is upper bounded by

$$O\left(\frac{\text{UB}}{\tilde{p}_{j^*}} \cdot (1/\varepsilon)\log(1/\varepsilon)\frac{\tilde{p}_{j^*}}{\varepsilon 2^\ell}\right) = O((1/\varepsilon^3)\log(1/\varepsilon)).$$

To show the claim, consider $\mathcal{S}'$ before entering Step 11 together with the corresponding set $\overline{\mathcal{M}}$ of machines that process some small job. Since $\mathcal{S}$ is the output of Online LPT, then the difference between the maximum an minimum loads of machines in $\overline{\mathcal{M}} \cap \mathcal{M}^{=}_{|\tilde{P}|}$ for solution $\mathcal{S}'$ is at most $2^\ell$. We call this property (P1). Also, notice that $\overline{\mathcal{M}} \cap \mathcal{M}^{\neq}_{|\tilde{P}|} = \emptyset$, and hence, the maximum load difference of two machines in this set is at most $2^\ell$, vacuously. We refer to this property as (P2). Notice that (P1) and (P2) hold iteratively throughout the later steps of the algorithm. Additionally, if some job is assigned to a machine $\mathcal{M}^{=}_{|\tilde{P}|}$ in Step 11, the algorithm does not enter the while loop and we

are done. Otherwise, the minimum load is achieved at $\mathcal{M}^{\neq}_{|\tilde{P}|}$. Hence, if there is a job migrated from a machine in $\mathcal{M}^{=}_{|\tilde{P}|}$ to $\mathcal{M}^{\neq}_{|\tilde{P}|}$ then the algorithm finishes. The claim follows. □

### 2.3 A note on geometric v/s arithmetic rounding

One of the main reasons to use our rounding procedure to multiples of $\varepsilon 2^{\ell}$ instead of the geometric rounding (i.e., down to the nearest power of $(1+\varepsilon)$) is because the same arguments used in this work cannot be applied to geometric rounded instances. It is crucial in the analysis that the number of possible loads is poly $(1/\varepsilon)$, while for geometric rounded instances that is not true as the following lemmas show.

**Lemma 17.** *Let $\varepsilon \in \mathbb{Q}_+, \varepsilon < 1$. Given a machine covering instance $(\mathcal{J}, \mathcal{M})$, let $\tilde{\mathcal{J}}$ be the set of jobs obtained by rounding geometrically jobs with processing time $p_j \in [\varepsilon \mathrm{OPT}, \mathrm{OPT}]$. If $C_1, C_2 \subseteq \tilde{J}$ are two different multi-sets of jobs with processing times at least $\varepsilon$OPT such that $\sum_{j \in C_i} p_j \in [\varepsilon \mathrm{OPT}, \mathrm{OPT}]$, $i = 1, 2$, then $\sum_{j \in C_1} p_j \neq \sum_{j \in C_2} p_j$.*

*Proof.* Assume w.l.o.g. OPT $= 1$. Hence the possible processing times are $(1+\varepsilon)^i$, with $i$ such that $\varepsilon \leq (1+\varepsilon)^i \leq 1$ (a finite family of such possible values). Suppose by contradiction that there are two different non-empty multi-sets $C_1, C_2$ with the same total load, and assume they are minimal, i.e. that there is no other pair of non-empty multi-sets with the same total load but with smaller total load. For $k = 1, 2$, let $C_k(j)$ be the number of jobs with processing time $(1+\varepsilon)^j$ in set $C_k$. Since the pair $C_1$, $C_2$ is minimal, we have that $C_1(j) = 0$ or $C_2(j) = 0$ for every $j$. $C_1$ and $C_2$ having the same total load means that

$$\sum_{j=-k}^{0} C_1(j)(1+\varepsilon)^j = \sum_{j=-k}^{0} C_2(j)(1+\varepsilon)^j,$$

where $k = -\lfloor \log_{1+\varepsilon}(\varepsilon) \rfloor$. This last equality can be rephrased as the existence of a non-zero polynomial $p(x) = b_0 + b_1 x + \cdots + b_k x^k$, with $|b_j| \in \{C_1(j), C_2(j)\}$ (i.e. with integer coefficients), that has $(1+\varepsilon)$ as one of its roots. Since $\varepsilon = \frac{c}{d} > 0$ for some co-primes $c$ and $d$, then $1 + \varepsilon = \frac{c+d}{d}$. Dividing $p(x)$ by $(dx - (c+d))$ leads to a polynomial $q(x) = a_0 + a_1 x + \cdots + a_{k-1} x^{k-1}$ which, thanks to Gauss lemma, has integer coefficients too. Let $b_i$ be the first coefficient of $p$ different from zero. Then

$$|b_i| = |(c+d)a_i + d a_{i-1}| = |\left(c + \frac{c}{\varepsilon}\right) a_i + \frac{c}{\varepsilon} a_{i-1}| > \frac{1}{\varepsilon},$$

implying, since the size of each job is at least $\varepsilon$, that the total load of the multi-sets is at least $b_i \varepsilon > 1$, which is a contradiction. □

**Lemma 18.** *Given $0 < \varepsilon < 1$, the number of different multi-sets of jobs with processing time at least $\varepsilon$OPT with total load at most* OPT *for a geometrically rounded instance is $2^{\Omega(\frac{1}{\varepsilon})}$.*

*Proof.* Let $u = \lfloor \log_2 \mathrm{OPT} \rfloor$ and $\ell = \lceil \log_2(\varepsilon \mathrm{OPT}) \rceil$. We will give a lower bound on the number of different sets with total load $2^u$ when the jobs are rounded to powers of 2, which implies that for $0 < \varepsilon < 1$ the same bound holds for processing times rounded to powers of $(1+\varepsilon)$. Let $\mathcal{C}_i$ be the number of different multi-sets with total load $2^{\ell+i}$. This number is characterized by the recurrence

$$C_0 = 1$$
$$C_{i+1} = 1 + \frac{C_i(C_i+1)}{2}.$$

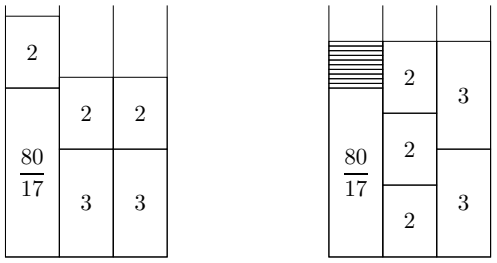


Figure 3: Left: Unique (17/16)-approximate solution before the arrival of small jobs. Right: Unique (17/16)-approximate solution after the arrival of small jobs.

This last term comes from the fact that a multi-set with total load $2^{\ell+i+1}$ can be constructed using only one job of size $2^{\ell+i+1}$, or merging two multi-sets of size $2^{\ell+i}$ (there are $\binom{C_i}{2} + C_i = \frac{C_i(C_i+1)}{2}$ such pairs).

Since recurrence $a_0 = 1$, $a_i = \frac{a_i^2}{2}$ satisfies $a_i \geq 2^{2^i}$, we conclude that $C_{u-\ell} \geq 2^{2^{\log \frac{1}{\varepsilon}}} \geq 2^{\Omega\left(\frac{1}{\varepsilon}\right)}$. □

Because of these two lemmas, if we use geometrically rounded instances we cannot make sure that, when a new jobs arrives to the system, the load profile changes only by poly $(1/\varepsilon)$ coordinates since there are $2^{\Omega(1/\varepsilon)}$ number of possible different loads.

## 2.4 An improved lower bound for the competitive ratio with constant migration factor

In opposition to the online makespan minimization with migration, where competitive ratio arbitrarily close to one can be achieved using constant migration factor [18], the online machine covering problem does not allow it. Until now, the best lower bound known for this ratio is $\frac{20}{19}$ [20], which we now improve to $\frac{17}{16}$ using similar ideas.

**Lemma 19.** *For any $\varepsilon > 0$, there is no $\left(\frac{17}{16} - \varepsilon\right)$-competitive algorithm using constant migration factor for the online machine covering problem with migration.*

*Proof.* Consider an instance consisting of 3 machines and 6 jobs of sizes $p_1 = p_2 = p_3 = 2$, $p_4 = p_5 = 3$ and $p_6 = \frac{80}{17}$. It is easy to see that the optimal solution is given by Figure 3 (a). Moreover, there is no other $\left(\frac{17}{16} - \varepsilon\right)$-approximate solution (up to symmetry).

Suppose by contradiction that there exists a $\left(\frac{17}{16} - \varepsilon\right)$-competitive algorithm with constant migration factor $C$. While processing the above instance, the algorithm must construct the optimal solution depicted in Figure 3 (left). Consider now that jobs with processing time smaller than $1/C$ arrive to the system, with total processing time $\frac{22}{17}$. Since the migration factor is $C$, none of the six previous jobs can be migrated, thus the best minimum load we can obtain is $\frac{96}{17}$, while the optimal solution is 6 (see Figure 3 (right)). We conclude by noting that $\frac{6}{96/17} = \frac{17}{16}$. □

Notice that the instance reaching the lower bound crucially depends on the arrival of jobs with arbitrarily small processing times. This kind of jobs are in fact the problematic ones, because under the assumption that at each iteration the incoming job is big enough (has processing time at least $\varepsilon$OPT), there is a robust PTAS with constant migration factor [20].

## 3 Jump and swap optimality for machine covering.

Let $\mathcal{F}$ denote the set of feasible solutions of a general combinatorial maximization problem with weight function $f\colon \mathcal{F} \to \mathbb{R}$. A *neighborhood function* is a mapping $N\colon \mathcal{F} \to 2^{\mathcal{F}}$ that specifies for each solution $s \in \mathcal{F}$ a set $N(s) \subseteq \mathcal{F}$, called *neighborhood* of $s$. The solutions in $N(s)$ are called *neighbors* of $s$, and a solution $\tilde{s} \in \mathcal{F}$ is called *locally optimal with respect to* $N$ or *$N$-optimal* if $f(\tilde{s}) \geq f(s)$ for every $s \in N(\tilde{s})$. A local move is a replacement of a solution by one of its neighbors. Any algorithm that starts from a feasible solution and iteratively perform local moves that improve the current solution until no such local move is possible is called a *local search algorithm*.

It is a good strategy to look for local search algorithms with good approximation guarantees whose locally optimal solution can be found efficiently, because they can be effectively adapted to the online context with migration. The analysis of their migration factor is done using the following premise: the migrated load corresponds to the sum of the migrated jobs in each local move, and for simplified instances (rounded, for example) the number of local moves until a locally optimal solution is found is usually a constant. That will be the case for two natural neighborhoods used in the load balancing context: *jump* and *swap*.

Two solutions are jump-neighbors if they assign the jobs to the same machines (up to symmetry) except for at most one job. A solution $\mathcal{S}'$ is a swap-neighbors of $\mathcal{S}$ if they assign the jobs to the same machines (up to symmetry) except for at most two jobs and, if they differ in exactly two jobs $j_1, j_2$ then they are in swapped machines, i.e., $\mathcal{S}(j_1) = \mathcal{S}'(j_2)$ and $\mathcal{S}(j_2) = \mathcal{S}(j_1)$. The *weight* of a solution is defined through a two-dimensional vector having the minimum load of the schedule as first coordinate and the number of non-least loaded machines as second ones. We compare the weight of two solutions lexicographically[1].

In other words, a solution is jump-optimal (respectively swap-optimal) if the migration of a single job (resp. the migration of a job or the swapping of two jobs) does not increase the minimum load and, if it maintains the minimum load then it does not reduce the number of least loaded machines. The following lemma characterizes jump-optimal solutions for machine covering.

**Lemma 20.** *Given $(\mathcal{J}, \mathcal{M})$ a machine covering instance, a schedule $\mathcal{S}$ is jump optimal if and only if for any machine $i \in \mathcal{M}$ and any job $j \in \mathcal{S}^{-1}(i)$, we have that $\ell_i(\mathcal{S}) - p_j \leq \ell_{\min}(\mathcal{S})$.*

*Proof.* If $\mathcal{S}$ is jump-optimal and there is a job not satisfying the inequality, then moving it to a least loaded machine either increases the minimum load of the schedule or reduces the number of least loaded machines, which is a contradiction.

On the other hand, if $\mathcal{S}$ is not jump-optimal then there is a job $j$ whose migration improves the weight of the solution or, if not, a job whose migration to a least loaded machine decreases the number of least loaded machines. Consider first the case in which moving $j$ from a machine $i$ increases the minimum load. This means that the new load of machine $i$, $\ell_i(\mathcal{S}) - p_j$, is at least the new minimum load, which is strictly larger than $\ell_{\min}(\mathcal{S})$, proving the needed inequality. Consider now the case in which moving $j$ from a machine $i$ to machine $i'$ maintains the minimum load while reducing the number of least loaded machines. Then $i'$ must have been a non-unique least loaded machine in $\mathcal{S}$. Furthermore, machine $i$ cannot become a minimum loaded machine in the new schedule (otherwise the number of least loaded machines would not change). This means that $\ell_i(\mathcal{S}) - p_j$ is strictly larger than $\ell_{\min}(\mathcal{S})$. □

[1] Just using the minimum load does not lead to good approximation ratios: think for example of $m > 2$ machines and $m$ jobs of size 1; it is jump-optimal to assign all of them to the same machine.

Now we show a simple upper bound for the approximation ratio of jump-optimality for machine covering.

**Lemma 21.** *Given an instance $(\mathcal{J}, \mathcal{M})$ for the machine covering problem, any jump-optimal solution $\mathcal{S}$ is $\frac{7}{4}$-approximate.*

*Proof.* For this proof we will not assume that the processing times are rounded. Observe first that every job $j$ with $p_j > \ell_{\min}(\mathcal{S})$ must be assigned in $\mathcal{S}$ to a machine $i$ on its own. If not, there would be another job $j'$ in machine $i$ such that $\ell(i) - j' \geq p_j > \ell_{\min}(\mathcal{S})$ which contradicts Lemma 20). Therefore, we can remove both the job $j$ and one machine from the instance. By doing this, the restricted schedules $\mathcal{S}$ is still jump-optimal, and the optimal value of the new instance cannot decrease: by arbitrarily reassigning the jobs that were in the same machine as $j$ in the optimal schedule to other machines we obtain a schedule for the new instance with the same or larger minimum load. As a consequence, we can assume by scaling the processing times that $\ell_{\min}(\mathcal{S}) = 1$, and that $p_j \leq 1$ for all $j \in \mathcal{J}$.

Let $\mathcal{M}_1$ be the set of machines $i \in \mathcal{M}$ such that $\mathcal{S}$ assigns exactly two jobs to $i$ and $\ell_i(\mathcal{S}) > 1$. Hence every machine $i$ outside $\mathcal{M}_1$ either has a load of 1, or it contains at least 3 jobs. In the latter case, we deduce there is a job $j$ with $p_j \leq \ell_i(\mathcal{S})/3$. By using Lemma 20 we conclude that all machines not in $\mathcal{M}_1$ have load at most 3/2. Consider now a modified instance $(\tilde{\mathcal{J}}, \mathcal{M})$ where all the jobs assigned to machines in $\mathcal{M}_1$ are replaced by jobs of size 1. $\mathcal{S}$ is still jump-optimal in the modified instance, it has the same minimum load and the optimal value cannot decrease because we are only enlarging the processing times. For this modified instance we have two cases:

1. If $|\mathcal{M}_1| \leq \frac{m}{2}$, then

$$\mathrm{OPT} \leq \tfrac{1}{m} \sum_{i=1}^{m} \ell_i(\mathcal{S}) \leq \frac{1}{m} \left( |\mathcal{M}_1| \cdot 2 + (m - |\mathcal{M}_1|) \frac{3}{2} \right) = \frac{1}{m} \left( \frac{3}{2} m + \frac{1}{2} |\mathcal{M}_1| \right) \leq \frac{7}{4}$$

2. If $|\mathcal{M}_1| > \frac{m}{2}$, then $\mathcal{S}_{\mathrm{OPT}}(\tilde{\mathcal{J}}, \mathcal{M})$ must assign to at least $2|\mathcal{M}_1| - m$ machines two jobs of size 1 and nothing else. This follows from the pigeonhole principle and the fact that no machine receives more than 2 jobs of size 1. Let $\mathcal{M}' \subseteq \mathcal{M}$ be such a set of $2|\mathcal{M}_1| - m$ machines and $\mathcal{J}' \subseteq \mathcal{J}$ the jobs assigned to $\mathcal{M}'$ in $\mathcal{S}_{\mathrm{OPT}}(\tilde{\mathcal{J}}, \mathcal{M})$. The instance $(\mathcal{J} \setminus \mathcal{J}', \mathcal{M} \setminus \mathcal{M}')$ together with the restriction of $\mathcal{S}$ to $\mathcal{J} \setminus \mathcal{J}'$ satisfies the hypothesis of the previous case, so by induction, $\mathrm{OPT}(\tilde{\mathcal{J}}, \mathcal{M}) \leq \mathrm{OPT}(\mathcal{J} \setminus \mathcal{J}', \mathcal{M} \setminus \mathcal{M}') \leq 7/4$.

Since all modifications cannot decrease the optimal value and do not change $\ell_{\min}(\mathcal{S})$, we conclude that $\mathcal{S}$ is a $\frac{7}{4}$-approximate solution for the machine covering problem. □

To get a lower bound in the approximation ratio $r$, we will use Sylvester's sequence $t := (2, 3, 7, 43, 1807, \dots)$, also defined by Salzer [17],

$$t_1 = 2; \qquad t_{k+1} = t_k(t_k - 1) + 1, \quad \text{for } k \geq 1, \tag{3}$$

and the related constant $h_\infty := \displaystyle\sum_{i=1}^{\infty} \frac{1}{t_i - 1} \approx 1.69103$, which appears many times in the design of natural approximation algorithms for packing problems [2, 4, 14]. First, we give some properties of the previous sequence which will be useful for the construction of the instances.

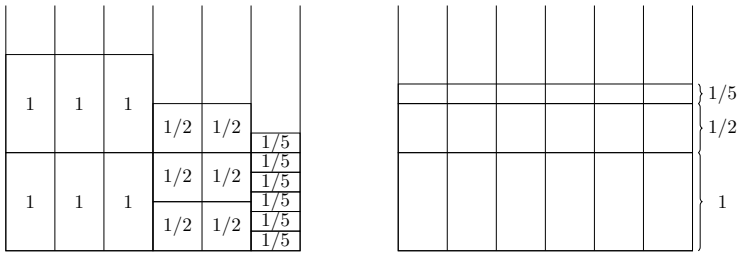


Figure 4: $\mathcal{S}_3^{\text{jump}}$ and the optimal solution for the respective instance $(\mathcal{J}_3^{\text{jump}}, \mathcal{M}_3^{\text{jump}})$.

**Lemma 22.** *For any $k \geq 2$, the sequence $t_k$ defined by* (3) *satisfies* $\left(\sum_{i=1}^{k-1} \frac{t_k - 1}{t_i}\right) + 1 = t_k - 1$, *and* $t_k = \left(\prod_{i=1}^{k-1} t_i\right) + 1$.

*Proof.* 1. We will prove the proposition by induction in $k$. The base case $k = 2$ is clearly true. Furthermore, for $k \geq 2$,

$$\begin{aligned}\sum_{i=1}^{k} \frac{t_{k+1} - 1}{t_i} + 1 &= \sum_{i=1}^{k} \frac{t_k(t_k - 1)}{t_i} + 1 = t_k \left(\frac{t_k - 1}{t_k} + \sum_{i=1}^{k-1} \frac{t_k - 1}{t_i}\right) + 1 \\ &= t_k \left(\frac{t_k - 1}{t_k} + t_k - 2\right) + 1 = t_k(t_k - 1) = t_{k+1} - 1.\end{aligned}$$

2. We will prove the equivalence term to term by induction in $k$. The case $k = 2$ is true by definition. Now, for $k \geq 2$

$$t_{k+1} = t_k(t_k - 1) + 1 = t_k \left(\prod_{i=1}^{k-1} t_i\right) + 1 = \left(\prod_{i=1}^{k} t_i\right) + 1.$$

□

Now, consider the following family of instances. For $k \geq 2$, let $\mathcal{M}_k^{\text{jump}}$ be a set of $t_k - 1$ machines, and $\mathcal{J}_k^{\text{jump}}$ a set with $t_k - 1$ jobs of size $(t_i - 1)^{-1}$ for each $i = 1, \ldots, k-1$, and $t_k - 1$ jobs of size $(t_k - 2)^{-1}$. Notice that $|\mathcal{J}_k^{\text{jump}}| = k(t_k - 1) = k|\mathcal{M}_k^{\text{jump}}|$. It is not difficult to see that, in the optimal solution, each machine will have assigned one job of each kind, implying that $\text{OPT}(\mathcal{J}_k^{\text{jump}}, \mathcal{M}_k^{\text{jump}}) = \left(\sum_{i=1}^{k-1} \frac{1}{t_i - 1}\right) + \frac{1}{t_k - 2}$.

The constructed solution $\mathcal{S}_k^{\text{jump}}$ assigns the jobs in the following way (see Figure 4 for an example): the $t_k - 1$ jobs of size $(t_k - 2)^{-1}$ are assigned to a single machine, and for $i = 1, \ldots, k-1$, the jobs of size $(t_i - 1)^{-1}$ are assigned to $(t_k - 1)/t_i$ machines, each one having $t_i$ jobs.

Lemma 22 implies that $(t_k - 1)/t_i$ is an integer and thus this assignment is valid, yielding a minimum load of $\ell_{\min}(\mathcal{S}_k) = 1 + (t_k - 2)^{-1}$. Also, $\mathcal{S}_k^{\text{jump}}$ is jump-optimal because migrating a job to the least loaded machine will leave the original machine with a load of 1, which would

decrease the minimum load. Hence, the approximation ratio $r$ must satisfy, for any $k \geq 2$, that

$$r \geq \frac{\left(\sum_{i=1}^{k-1} \frac{1}{t_i - 1}\right) + \frac{1}{t_k - 2}}{1 + \frac{1}{t_k - 2}}.$$

Taking $k \to \infty$ and together with Lemma 21, we conclude the following result.

**Theorem 23.** *For the machine covering problem, any jump-optimal schedule $\mathcal{S}$ is $r$-approximate, where $r \in [h_\infty, \frac{7}{4}] \approx [1.691, 1.75]$.*

**Remark.** *The lower and upper bound also apply for the approximation ratio of swap-optimality: the schedules $\mathcal{S}_k^{jump}$ are swap-optimal, and any swap-optimal schedule is also jump-optimal by definition.*

### 3.1 From local search to bounded migration.

**Algorithm 2** Push

**Input:** Schedule $\mathcal{S}$ for $(\mathcal{J}, \mathcal{M})$, $i \in \mathcal{M}$, $j \notin \mathcal{J}$
**Output:** $Q \subseteq \mathcal{J}$, schedule $\mathcal{S}'$ for $((\mathcal{J} \cup \{j\}) \setminus Q, \mathcal{M})$

```
Q ← ∅.
S' ← S.
assign j to machine i in S'.
for k ∈ S^{-1}(i) do
    if ℓ_i(S') − p̃_k > ℓ_min(S') then
        take out k from i in S'.
        Q ← Q ∪ {k}.
    end if
end for
return Q, S'.
```

Using the rounding procedure developed in Section 2.2, jump-optimality can be adapted to the online context using migration factor $O\left(\frac{1}{\varepsilon}\right)$. In order to do this, we need an auxiliary algorithm called *Push* (Algorithm 2) to assign a job $j$ to a given machine. The algorithm assigns the job to the machine and then iteratively push out from that machine all jobs breaking jump-optimality, storing them in a special set which is part of the output. This algorithm is the base of the push neighborhood analyzed by Schuurman and Vredeveld [19].

Our algorithm, which is described in detail in Algorithm 3, is called every time a new job $j^*$ arrives to the system, and will receive as input the current solution $\mathcal{S}$ for $(\mathcal{J}, \mathcal{M})$, initialized as empty if $\mathcal{J} = \emptyset$. It will output a 4-relaxed version of a jump-optimal solution. The general idea of Algorithm 3 is to first round the instance, and then use the Push algorithm to assign the incoming job to a least loaded machine. This can push out some jobs that need to be reassigned. We assign them iteratively using Push until a jump-optimal solution is obtained. Notice that assigning and pushing jobs can be interpreted as cleverly moving from neighbor to neighbor in the jump neighborhood. As a last step, all small jobs that were pushed during the entire procedure are reassigned using list-scheduling.

Similarly to Algorithm 1, this procedure is called every time a new job $j^*$ arrives to the system, and will receive as input the current solution $\mathcal{S}$ for $(\mathcal{J}, \mathcal{M})$, initialized as empty if $\mathcal{J} = \emptyset$. In

the algorithm, $\mathcal{S}'_B$ denotes schedule $\mathcal{S}'$ restricted to big jobs. For simplicity, we assume that every arriving job is either small or big. It is easy to extend the algorithm to handle huge jobs.

**Algorithm 3** Online jump-optimality

**Input:** Instances $(\mathcal{J}, \mathcal{M})$ and $(\mathcal{J}', \mathcal{M})$ such that $\mathcal{J}' = \mathcal{J} \cup \{j^*\}$; a schedule $\mathcal{S}(\mathcal{J}, \mathcal{M})$.

run LPT on input $\mathcal{J}'$ and let $\tau$ be the minimum load. Set UB $\leftarrow 2\tau$. Define $\tilde{P}, \ell$, and $u$ based on this upper bound UB using (1) and (2).
set $\mathcal{S}' \leftarrow \mathcal{S}$
**if** $\tilde{p}_{j^*} < 2^\ell$ **then**.
    assign $j^*$ to a least loaded machine in $\mathcal{S}'$.
**else**
    set $Q_B \leftarrow \{j^*\}$. ▷ `Set with unassigned big jobs.`
    set $Q_s \leftarrow \emptyset$. ▷ `Set with unassigned small jobs.`
    **while** $Q_B \neq \emptyset$ **do**
        for any $j \in Q_B$, $Q_B \leftarrow Q_B \setminus \{j\}$.
        in $\mathcal{S}'_B$, use Algorithm 2 to assign $j$ to a least loaded machine $m^*$, obtaining its output set $Q$. Update $\mathcal{S}'_B$ to be the output solution of this procedure.
        reassign jobs in $\mathcal{S}'$ such that the assignment of (big) jobs in $\mathcal{S}'$ and $\mathcal{S}'_B$ coincides.
        **while** $m^*$ contains a small job w.r.t. UB and $\ell_{m^*}(\mathcal{S}') > \ell_{\min}(\mathcal{S}') + 2^\ell$ **do**
            remove the smallest job in ${\mathcal{S}'}^{-1}(m^*)$ and add it to $Q_s$.
        **end while**
        $Q_B \leftarrow Q_B \cup Q$.
    **end while**
    assign the jobs in $Q_s$ to $\mathcal{S}'$ using list-scheduling.
**end if**
**return** $\mathcal{S}'$.

Using similar techniques as in Lemma 11, we can argue about the approximation ratio of the algorithm.

**Lemma 24.** *Algorithm 3 is $(r+\varepsilon)$-competitive and has polynomial running time.*

*Proof.* In order to show that the competitive ratio is $(r+\varepsilon)$, we just need to show that the schedule constructed by Algorithm 3 is a $k$-relaxed version of a jump-optimal solution for some convenient $k$. This follows from Theorem 23 and Lemma 3.

First, Algorithm 3 has, for any $h \in \mathbb{R}^+$, robust structure to level $h$ because each time that Push is called, it moves a total load of jobs smaller than the processing time of the inserted job (otherwise, the machine would not be a least loaded machine), and if the job is small, then nothing is migrated. This also directly implies that the running time of the algorithm is polynomial because every job is migrated at most once, so each while loop is executed only a polynomial number of times.

Now, let us show that the constructed schedule is a 4-relaxed version of a jump-optimal schedule. Let $k = \frac{2^\ell}{\varepsilon \mathrm{OPT}'}$, which implies that big jobs have processing time at least $k\varepsilon\mathrm{OPT}'$. If $\tilde{p}_{j^*} < 2^\ell$, it is easy to see that the conditions are fulfilled. If job $j^*$ is big, we first notice that in each iteration of the *while* loop indexed by $Q_B$, after applying Push and updating, the resulting schedule restricted to big jobs $\mathcal{S}'_B$ is jump-optimal if the jobs in $Q_B$ are excluded. When this *while* loop is finished, since

$Q_B = \emptyset$, the partial schedule is jump-optimal restricted to big jobs, fulfilling the first condition of being $k$-relaxed. Also, the internal *while* loop makes sure that at each iteration the second condition of being $k$-relaxed is satisfied, so at the end of the *while* loop indexed by $Q_B$, the partial schedule is a $k$-relaxed version of a jump-optimal solution. By assigning the remaining small jobs using list-scheduling, the final solution is a $k$-relaxed version of a jump-optimal solution, and we finish by noticing as in Lemma 11 that $k \leq 4$. □

To analyze the migration factor, we define the *migration tree* of the algorithm as a node-weighted tree $G = (V, E)$, where $V$ is the set of migrated jobs together with the incoming job $j^* \notin \mathcal{J}$. The weight of each node $v \in V$ is the processing time of the corresponding job $\tilde{p}_v$. The tree is constructed by first adding job $j^*$ as root. For each node (job) $v$ in the tree that is migrated, its children are defined as all the jobs migrated at the insertion of $v$. It is easy to see that this process does not create any loops as the processing time of a child is strictly smaller than the processing time of its parent. By definition the leaves of the tree are the jobs not inducing migration, and thus any small job in the tree is a leaf. In the context of local search, the number of nodes in the tree corresponds to the number of iterations of the specific local search procedure. The following lemma bounds the migration factor of Algorithm 3.

**Lemma 25.** *Algorithm 3 has migration factor* $O(1/\varepsilon)$.

*Proof.* Let $w_i$ be the total processing time of nodes corresponding to big jobs in level $i$ of the migration tree. Assume that $\tilde{p}_{j^*} = q_\kappa = 2^g + h\varepsilon 2^g$ for some $\kappa \in \{1, \ldots, |\tilde{P}|\}$, $g \in \{\ell, \ldots, u\}$ and $h \in \{0, \ldots, \frac{1}{\varepsilon} - 1\}$. Every time a job $j$ is inserted using Push, the total load of jobs in the output $Q$ of the algorithm is strictly less than $\tilde{p}_j$, which means that $w_i$ is strictly decreasing, and also that at each level $i$ of the tree there are at most $\frac{w_i}{2^\ell}$ nodes corresponding to big jobs. Since the second condition of being a $k$-relaxed version (Definition 2) of a jump-optimal solution is maintained through the iterations, the small jobs that need to be migrated because of insertion of a big job $j$ have total load at most $\tilde{p}_j + 2^\ell$. This implies that the total load of small jobs at each level $i \geq 1$ of the tree is at most $w_{i-1} + \frac{w_{i-1}}{2^\ell} \cdot 2^\ell = 2w_{i-1}$, and hence the total processing time of nodes corresponding to small jobs is at most twice the total processing time of nodes corresponding to big jobs.

From now on we assume that the migration tree contains only nodes corresponding to big jobs. We categorize each level $i \geq 1$ of the migration tree according to the following two cases: if there is a node in level $i-1$ having at least two children, we say that level $i$ falls in case 1, and it falls in case 2 otherwise. We first show that there are at most $\frac{\tilde{p}_{j^*}}{2^\ell} \leq 1/\varepsilon$ levels of the tree falling in case 1. Because of the way the migration tree is constructed, it is not difficult to see that the total weight of the leaves in the tree is at most $\tilde{p}_{j^*}$ (this property is maintained inductively through the executions of Algorithm Push). Because of this, since each big job has processing time at least $2^\ell$, every migration tree has at most $\tilde{p}_{j^*}/2^\ell$ leaves, which is also an upper bound for the number of nodes that have more than one children in the tree (each one of them induces at least one extra leaf), and hence for the number of levels falling in case 1.

There can be more than $1/\varepsilon$ levels falling in case 2 along the tree, but we will show that in that case $w_i$ quickly decreases based on the following claim.

**Claim:** Let $q_{i_1}, \ldots, q_{i_k} \in \tilde{P}$ such that $\sum_{j=1}^{k} q_{i_j} \in (q_{s+1}, q_s]$ for some $s \in \{1, \ldots, |\tilde{P}|\}$. Then $\sum_{j=1}^{k} q_{i_j+1} \leq$

$\sum_{j=1}^{k} q_{i_j} - \frac{\varepsilon}{4} q_{s+1}$, where we assume that $q_{|\tilde{P}|+1} = 0$.

Notice that the claim implies that for a level $i$ falling in case 2, if $w_{i-1} \in (q_{s+1}, q_s]$ for some $s \in \{1, \ldots, |\tilde{P}|\}$, then $w_i \leq w_{i-1} - \frac{\varepsilon}{4} q_{s+1}$. To compute the total processing time of the nodes in the migration tree, we will bound the total weight of the levels corresponding to each case separately. Since there are at most $1/\varepsilon$ levels falling in case 1, each one of them having total weight at most $\tilde{p}_{j^*}$, we can bound the total weight of those levels by $\frac{1}{\varepsilon}\tilde{p}_{j^*}$. Let us now relabel the levels of the tree where the second case occurs by just $\{1, 2, \ldots, L_2\}$ (i.e. we ignore the levels falling in case 1). Thanks to the claim, for every $i \in \{1, 2, \ldots, L_2\}$, if $w_{i-1} \in (q_{s+1}, q_s]$ for some $s \in \{1, \ldots, |\tilde{P}|\}$, then $\sum_{j=i-1}^{i+2} w_j \leq 4q_s$ and $w_{i+3} \leq q_{s+1}$ (because $q_s - \varepsilon q_{s+1} \leq q_{s+1}$), and we can restart the process for $i+3$ with the correct value $q_{s'} \leq q_{s+1}$. If we use this argument starting with $w_0 \in (q_{\kappa+1}, q_\kappa]$, we can conclude that $\sum_{i=0}^{L_2} w_i \leq 4 \sum_{i=\kappa}^{|\tilde{P}|} q_i$, which, recalling that $\tilde{p}_{j^*} = 2^g + h\varepsilon 2^g$, is at most

$$4 \sum_{i=\ell}^{g} \sum_{b=0}^{\frac{1}{\varepsilon}-1} (2^i + b\varepsilon 2^i) = 4 \sum_{i=\ell}^{g} 2^{i-1} \left(\frac{3}{\varepsilon} - 1\right) \leq 4\tilde{p}_{j^*} \left(\frac{3}{\varepsilon} - 1\right).$$

These two bounds, together with the fact that the total load of small migrated jobs is at most twice this value, implies that the migration factor is at most $O\left(\frac{1}{\varepsilon}\right)$.

To prove the claim, notice first that

$$\sum_{j=1}^{k} q_{i_j+1} \leq \sum_{j=1}^{k} q_{i_j} - \sum_{j=1}^{k} \varepsilon 2^{\lfloor \log(q_{i_j}) \rfloor - 1} \leq \sum_{j=1}^{k} q_{i_j} - \frac{\varepsilon}{2} \sum_{j=1}^{k} 2^{\lfloor \log(q_{i_j}) \rfloor}.$$

Also, since $\sum_{j=1}^{k} 2^{\lceil \log(q_{i_j}) \rceil} \geq \sum_{j=1}^{k} q_{i_j} > q_{s+1}$, we have that $\sum_{j=1}^{k} 2^{\lfloor \log(q_{i_j}) \rfloor} > \frac{q_{s+1}}{2}$. This concludes the proof of the claim. □

Moreover, the following lemma shows that the previous analysis is essentially tight.

**Lemma 26.** *There are instances for which Algorithm 3 uses a migration factor of at least* $\Omega\left(\frac{1}{\varepsilon}\right)$.

*Proof.* Consider an instance with $\text{OPT} = 2^{u+1}$ and $\varepsilon\text{OPT} = 2^\ell$ for some integers $\ell, u$, and assume for simplicity that UB = OPT. This way, $\tilde{I}(\text{UB}) = \{\ell, \ldots, u\}$. The instance, consisting of $m \in O\left(\frac{1}{\varepsilon} \log \frac{1}{\varepsilon}\right)$ machines, is constructed in the following way: Consider the possible processing times sorted non increasingly $t_1, \ldots, t_h$. For each $i$ such that $t_i < 2^u$, the schedule has a machine with a job of size $t_i$ assigned, and it is completed with jobs until having load $2^{u+1}$: if $t_i = 2^k + j\varepsilon 2^k$, this can be done adding a job of size $2^k + \left(\frac{1}{\varepsilon} - j\right)\varepsilon 2^k$, a job of size $2^k$ and for each $k' = k+2, \ldots, u$, a job of size $2^{k'}$ (if $i = u-1$, the machine will not have any of these last jobs). By doing so, the load of the machine is

$$2^i + j\varepsilon 2^i + 2^i + \left(\frac{1}{\varepsilon} - k\right)\varepsilon 2^i + 2^i + \sum_{i'=i+2}^{u} 2^{i'} = 2^{i+2} + 2^{u+1} - 2^{i+2} = 2^{u+1}.$$

Now, if a job of size $2^u$ arrives to the system, it can be inserted using Push in the machine with the largest job of size less than $2^u$ (i.e., with processing time $2^{u-1} + \left(\frac{1}{\varepsilon} - 1\right) \varepsilon 2^{u-1}$), taking out such job because it breaks jump-optimality. If Algorithm 3 takes the decision in the same way iteratively, then at least one job of each possible size $t_i < 2^u$ is migrated, being then the total migrated load at least

$$\sum_{i=\ell}^{u-1} \sum_{j=1}^{\frac{1}{\varepsilon}-1} 2^i + j\varepsilon 2^i = \left(\frac{1}{\varepsilon} - 1\right)(2^u - 2^{\ell+1}) + \frac{1}{2}\left(\frac{1}{\varepsilon} - 1\right) 2^u \in \Omega\left(\frac{1}{\varepsilon} 2^u\right),$$

and hence the migration factor needed for this instance is $\Omega\left(\frac{1}{\varepsilon}\right)$. □

Putting Lemmas 24, 25 and 26 together, we can conclude the following result.

**Theorem 27.** *Given $\varepsilon > 0$, Algorithm 3 is a polynomial time $(r + \varepsilon)$-competitive algorithm, where $r$ is the approximation ratio of jump-optimality, and uses migration factor $O(1/\varepsilon)$. Moreover, there are instances for which this factor is $\Omega(1/\varepsilon)$.*

# References


[1] Y. Azar and L. Epstein. On-line machine covering. *J. Sched.*, 1:67–77, 1998.

[2] N. Bansal, X. Han, K. Iwama, M. Sviridenko, and G. Zhang. Harmonic algorithm for 3-dimensional strip packing problem. In *SODA 2007*, pages 1197–1206.

[3] S. Berndt, K. Jansen, and K. Klein. Fully dynamic bin packing revisited. In *APPROX/RANDOM 2015*, pages 135–151.

[4] A. Caprara. Packing 2-dimensional bins in harmony. In *FOCS 2002*, pages 490–499.

[5] J. Csirik, H. Kellerer, and G. Woeginger. The exact LPT-bound for maximizing the minimum completion time. *Oper. Res. Lett.*, 11:281–287, 1992.

[6] B. Deuermeyer, D. Friesen, and M. Langston. Scheduling to maximize the minimum processor finish time in a multiprocessor system. *SIJADM*, 3:190–196, 1982.

[7] L. Epstein and A. Levin. A robust APTAS for the classical bin packing problem. *Math. Program.*, 119:33–49, 2009.

[8] L. Epstein and A. Levin. Robust algorithms for preemptive scheduling. *Algorithmica*, 69:26–57, 2014.

[9] A. Frangioni, E. Necciari, and M. Scutellà. A multi-exchange neighborhood for minimum makespan parallel machine scheduling problems. *J. Comb. Optim.*, 8:195–220, 2004.

[10] A. Gu, A. Gupta, and A. Kumar. The power of deferral: Maintaining a constant-competitive steiner tree online. *SIAM J. Comput.*, 45:1–28, 2016.

[11] K. Jansen and K. Klein. A robust AFPTAS for online bin packing with polynomial migration. In *ICALP 2013*, pages 589–600.

[12] K. Jansen, K. Klein, and J. Verschae. Closing the gap for makespan scheduling via sparsification techniques. In *ICALP 2016*, pages 1–13.

[13] J. Łacki, J. Oćwieja, M. Pilipczuk, P. Sankowski, and A. Zych. The power of dynamic distance oracles: Efficient dynamic algorithms for the steiner tree. In *STOC 2015*, pages 11–20.

[14] C. Lee and D. Lee. A simple on-line bin-packing algorithm. *J. ACM*, 32:562–572, 1985.

[15] N. Megow, M. Skutella, J. Verschae, and A. Wiese. The power of recourse for online MST and TSP. *SIAM J. Comput.*, 45:859–880, 2016.

[16] D. Recalde, C. Rutten, P. Schuurman, and T. Vredeveld. Local Search Performance Guarantees for Restricted Related Parallel Machine Scheduling. *LATIN 2010*, pages 108–119, 2010.

[17] H. Salzer. The approximation of numbers as sums of reciprocals. *Am. Math. Monthly*, 54:135–142, 1947.

[18] P. Sanders, N. Sivadasan, and M. Skutella. Online scheduling with bounded migration. *Math. Oper. Res.*, 34:481–498, 2009.

[19] P. Schuurman and T. Vredeveld. Performance guarantees of local search for multiprocessor scheduling. *INFORMS J. Comput.*, 19:52–63, 2007.

[20] M. Skutella and J. Verschae. Robust polynomial-time approximation schemes for parallel machine scheduling with job arrivals and departures. *Math. Oper. Res.*, 41:991–1021, 2016.

[21] G. Woeginger. A polynomial-time approximation scheme for maximizing the minimum machine completion time. *Oper. Res. Lett.*, 20:149–154, 1997.

# A Non-constant Migration for classic LPT

**Lemma 28.** *For any $k \geq 2$ there exists a set $\mathcal{J}$ of $4k+1$ jobs and an extra job $j^* \notin \mathcal{J}$ such that, for every schedule $\mathcal{S}$ constructed using* LPT *on $2k+1$ machines, it is not possible to construct a schedule $\mathcal{S}'$ using* LPT *for $\mathcal{J} \cup \{j^*\}$ with migration factor less than $m/2$.*

*Proof.* Fix a constant $0 < \varepsilon \leq \frac{1}{6k}$. Consider a set $\mathcal{J}$ consisting of the following $4k+1$ jobs: $k+1$ jobs of size 1; for each $i \in \{0, \ldots, k-1\}$, a job of size $\frac{1}{2}+i\varepsilon$ and a job of size $\frac{1}{2}-(i+1)\varepsilon$, and finally $k$ jobs of size $\frac{1}{2}-k\varepsilon \geq \frac{1}{3}$. Assume the jobs in $\mathcal{J}$ are sorted non-increasingly by size. There is a unique schedule constructed using LPT for this instance (up to symmetry) which assigns the jobs in the following way (see Figure 1a): The $k+1$ jobs of size 1 to a machine on their own, and for each $i = 1, \ldots, k$, it assigns to machine $k+i$ a job of size $\frac{1}{2}+(k-i-1)\varepsilon$, a job of size $\frac{1}{2}-(k-i)\varepsilon$ and a job of size $\frac{1}{2}-k\varepsilon$ (since the total load of the first two jobs is $1-\varepsilon$, the last $k$ jobs must be assigned to these $k$ machines).

Now consider an arriving job $j^*$ of size $\frac{1}{2}+k\varepsilon \leq \frac{2}{3}$. There is a unique schedule constructed using LPT for the new instance (up to symmetry) which assigns the jobs in the following way (see Figure 1b): it assigns to the first $k+1$ machines a job of size 1 and a job of size $\frac{1}{2}-k\varepsilon$, to machine $k+2$ job $p_{j^*}$ and a job of size $\frac{1}{2}-(k-1)\varepsilon$, for each $i = 2, \ldots, k-1$ it assigns to machine $k+i+1$ a job of size $\frac{1}{2}+(k+1-i)\varepsilon$ and a job of size $\frac{1}{2}-(k-i)\varepsilon$, and finally to machine $2k+1$ a job of size $\frac{1}{2}+\varepsilon$ and a job of size $\frac{1}{2}$ (now the total load of machines $k+2, \ldots, 2k+1$ is $1+\varepsilon$, then the last $k+1$ jobs must be assigned to the first $k+1$ jobs).

It is not difficult to see that, in the new schedule, every machine has a different subset of jobs assigned to it compared with the original schedule, and so at least one job must have been migrated per machine. Thus, the migrated total load is at least the load of the smallest $2k+1$ jobs, which implies that the needed migration factor is at least

$$\frac{\sum_{i=0}^{2k} p_{(4k+1)-i}}{p_{j^*}} \geq \frac{(2k+1)\left(\frac{1}{2}-k\varepsilon\right)}{\frac{1}{2}+k\varepsilon} \geq m\frac{1/3}{2/3} = \frac{m}{2}. \square$$